IMPERATIVE STATICS AND DYNAMICS

Nate Charlow (ncharlo@umich.edu)
Department of Philosophy
University of Michigan

CONTENTS

1 INTRODUCTION

Imperatives are linguistic devices used by an authority (speaker) to express wishes, requests, commands, orders, instructions, and suggestions to a subject (addressee).[1] This essay's goal is to tentatively address some of the following questions about the imperative.

- **Imperative Metasemantics.** What is the menu of options for understanding fundamental semantic notions like satisfaction, truth-conditions, validity, and entailment in the context of imperatives? Are there good imperative *arguments*, and, if so, how are they to be characterized? What are the options for understanding the property that an account of good imperative arguments is supposed to track? What constraints on a semantic analysis of the imperative do different positions on the metasemantic issues impose?

- **Imperative Semantics.** How might we implement metasemantic postures in a rigorous formal system? How much can we do using familiar tools from deontic modal logic? How much leverage over semantic questions can we gain by introducing tools from natural language semantics—ordering sources, dyadic modal

1. This is a crude conception of the function of the imperative in natural language, but it will do for our rather limited purposes. A more sophisticated treatment may be found in Hamblin (1987).

operators, salient alternatives, and the like—into a formal semantics for an imperative object language? How much leverage can we gain by introducing tools from rather less-utilized areas of modal logic—devices for representing actions and planning in time, modal operators constructed from action-terms, and the like—into the analysis?

- **Imperative Dynamics.** How do imperatives succeed in performing the speech-acts they are used to perform? How do imperatives update discourses? How can we leverage an account of imperative discourse update in giving a dynamic semantics for the imperative? Is there anything about the imperative that *demands* a dynamic semantic treatment?

Before jumping in, a very quick overview of the structure of the essay and general approach we will take to answering the above questions. In §2, we address metasemantic questions about the imperative, outlining in some detail three different tacks one might take in response to them. Very roughly: *fulfillment-oriented* logics regard the fundamental semantic relation—the imperative analogue of satisfaction—as fulfillment of the requirement expressed by an imperative. *Content-oriented* logics treat the fundamental semantic relation as requirement in view of an authority's desires. And *planning-oriented* logics treat the fundamental semantic relation as requirement in view of "prior" constraints on an agent's practical reasoning and intention formation. Each tack is naturally adopted to an idiosyncratic conception of validity in imperative argument. Depending on implementation, taking one tack over another naturally manifests in different predictions about the validity of certain argument patterns.

The rest of the essay is devoted to brass tacks—implementing different positions on the metasemantic questions in a formal semantics (both static and dynamic) for a formal imperative language. In §3, we develop a logic of content in terms of a slightly embellished, but otherwise fairly standard, deontic modal logic. We show that this setup is well-suited to handling an array of phenomena about the imperative—conditional, qualified, or otherwise hedged imperatives and the Ross Paradox, in particular—and we devote a significant amount of effort to exploring the intricacies surrounding each of these issues. The analysis of conditional imperatives is argued to benefit from the introduction of dyadic modal operators (and corresponding dyadic imperative operators) with restriction arguments and a novel version of an ordering-source semantics, rather than a simple accessibility relation semantics. The Ross Paradox is argued to benefit from the use of alternative semantics (although a Montague-Scott approach is also considered and rejected). We show that the setup can, with a rather minimal conceptual sleight of hand, avoid the trap of construing imperative operators in the imperative language as *literally* deontic modal operators, and thereby sidestep the problem of assigning formulas of the imperative language literal truth-conditions. Although the analysis is at bottom a treatment of an imperative logic of content in terms of the model theory for deontic modal logic, it is neither eliminativist nor reductionist.

In §4, we shift our attention to modeling how imperatives constrain the planning behavior of their subjects *at a given point in time*. Some of what we want to model in this area can be handled with a rather simple change to the apparatus of §3: replacing the authority-oriented ordering-source with a subject-oriented ordering-source—a *To-Do List* for an agent, roughly in the sense of Portner (2004, 2008)—although the shift introduces several complexities that are without direct parallel in the logic of content. Other theoretical goals demand a genuine elaboration of the apparatus. In shifting our focus to planning, it is natural—and, as we see, extremely useful—to have at our disposal a dynamic language that is capable of talking explicitly about actions and the hypothetical impact of the performance of certain actions on the constraints impinging on the agent's future planning. We adopt such a language—an embellishment of the Propositional Dynamic Language often used in providing a logic of programs—and proceed to rewrite the semantics from the ground up. The setup is shown to be extendable to a resolution of arbitrarily complex versions of the Ross Paradox, somewhat along the lines of that given in the prior section, as well as a perspicuous analysis of temporal imperative constructions and temporal constraints on planning that present difficulties when the imperative object language is non-dynamic. We close by suggesting a perspective—the perspective of "constraint semantics" in the sense of Swanson (2006, 2008a)—from which the logic of planning presented in this section could just as well be construed as a logic of content.

In §5, we turn to genuine dynamics, in particular modeling imperative effects on discourse. We do several things in this section. First, we introduce a new discourse parameter—a Rights List—to keep track of permissions, entitlements, and freedoms accumulated by an agent over the course of a discourse, and we explain the sort of work it can be used to do. Second, we define a set of update potentials—functions mapping from discourses into updated discourses—for the enriched, dynamic imperative object language. The update potentials are designed to make good on a particular conception of imperative force in discourse, on which imperatives both introduce constraints and grant rights with respect to the agent's future planning behavior. The result is an account of imperative effects on discourse that (i) handles a *much* wider array of imperative constructions than tackled in recent literature on the subject and (ii) offers a rather more subtle treatment of imperative force than what is offered in contemporary accounts. We use these definitions of update potentials to define irreducibly dynamic analogues of imperative satisfaction and entailment. We close with a defense of static semantics of the imperative in terms of the model theory for deontic modal logic against attempts to assert dynamic hegemony. The thesis we advance here can be read as the implicit and overarching theme of this paper: much of what we want to model about the imperative can be modeled using just the well-understood model theory for modal languages.

2 FOUNDATIONS

2.1 Motivating Imperative Logic

There is something undeniably compelling about the following argument patterns.

(1) Brush your teeth and go to bed!
$\therefore$ Brush your teeth!

(2) Pet every kitty!
$\therefore$ Pet Fluffy!

(3) If you go to Harlem, take the A-train!
Go to Harlem!
$\therefore$ Take the A-train!

(4) Use an axe or a saw!
Don't use an axe!
$\therefore$ Use a saw!

What precisely *makes* these argument patterns compelling—indeed, what it even means to describe the inference of an imperative conclusion from a set of imperative premises as compelling—is for now a bit of a mystery—one which will occupy our attention for much of the first part of this essay. But no matter. The pull which these arguments exert on us constitutes prima facie reason for thinking that there is a such a thing as *good im-*

perative arguments—and, therefore, a difference between good imperative arguments and bad imperative arguments.

It is the task of the imperative logician to characterize this difference, both semantically and syntactically. Semantically by attempting to define an imperative analogue of semantic validity or truth-preservation (and perhaps, although not necessarily, an imperative analogue of the model theories used to interpret non-imperative formal languages). Syntactically by attempting to develop a proof theory for a logical imperative object language. In this essay, we focus our attention on the semantic dimension of the project, and leave the proof theory for another occasion.

My own interest in imperatives stems from an interest in modeling the role of natural language imperatives in communication and individual/group practical reasoning. But, though we devote a large amount of time to developing a model for updating discourses and cognitive states in accordance with imperative utterances, our focus in this essay will not be on natural language—not directly anyway. Rather, we take the advice of Bar-Hillel (1966) and confine our attention to defining, interpreting, and developing a rough semantics and pragmatics for a formal imperative object language (with the customary expectation that insights from our intuitive understanding of the behavior and function of natural language imperatives will inform the formalism, and that a carefully designed and implemented formalism will enrich our understanding of imperative discourse).

For most of our purposes in this essay, an expressively poor imperative language built on top of the Boolean propositional language suffices. Let $\mathcal{L}_P$ be the Boolean propositional language. Then the Basic Propositional Imperative Language $\mathcal{L}_{PI}$ is defined as the smallest set such that:

(5) If $\phi \in \mathcal{L}_P$, then $\phi \in \mathcal{L}_{PI}$
If $\phi \in \mathcal{L}_P$, then $!\phi \in \mathcal{L}_{PI}$

The imperative "operator" $\ulcorner ! \urcorner$ may be read as "see to it that." Note that in limiting ourselves to a propositional language, we prevent ourselves from saying anything about the argument in form in (2). Nevertheless, there are, as we will see, plenty of interesting imperative arguments we can represent in $\mathcal{L}_{PI}$ and, once we have a semantics on the table, evaluate for validity. The following are natural $\mathcal{L}_{PI}$ logical forms for the arguments in (1), (3), and (4):

(6) $!(\phi \wedge \psi)\ /\ !\phi$

(7) $!(\phi \rightarrow \psi), !\phi\ /\ !\psi$

(8) $!(\phi \vee \psi), !\neg\phi\ /\ !\psi$

A note on $\mathcal{L}_{PI}$: it is natural (and customary) not to allow an imperative operator to take scope over another imperative operator, and we follow that custom here. More generally, we prohibit any embeddings of expressions containing the imperative operator. In view of the following sorts of constructions, this might be thought an intolerable expressive limitation on our language.

(9) If he doesn't stop, shoot!

(10) Stop, or I'll shoot!

A natural first pass at *conditional imperatives* like (9)—and, for that matter, the first "premise" imperative of (3)—and imperative threats like (10) would involve representing them with the logical forms $\neg\phi \rightarrow !\psi$ and $!\phi \vee \psi$ respectively. Such logical forms do not count as well-formed formulae of $\mathcal{L}_{PI}$. We will take up this issue again in our discussion of conditional imperatives, but a small piece of motivation for this decision might be in order. Sanctioning such logical forms seems to mean either interpreting imperative formulae as bearers of truth-values or devising non-truth-functional interpretations of their main Boolean connectives. Both options seem less than ideal. We avoid making the choice by provisionally supposing that these constructions are to be handled in some other way.[2]

With a basic imperative language in place, we now consider different possibilities for characterizing the property that distinguishes good imperative arguments from bad—viz., different proposals for understanding the imperative analogue of the entailment relation $\vDash$. Note that the menu of possible options is more diverse than represented here. Those presented here have been selected in view of their connection to the main themes of this essay (and interests of the author).

2.2 Logics of Fulfillment

The earliest proposal for conditions on imperative entailment was formulated in Jørgensen (1937-8), which proposed a reduction of imperative to "ordinary" (i.e., non-imperative) logic, along the following lines. (Note: we reserve $\Vdash$ for the imperative analogue of the entailment relation $\vDash$.)

(11) $!\phi_1, ..., !\phi_n \Vdash !\psi$ iff $\phi_1, ..., \phi_n \vDash \psi$

This analysis predicts the argument patterns in (1), (3), and (4) valid, on their suggested $\mathcal{L}_{PI}$ renderings. References that, with qualifications (to be discussed shortly), defend this basic picture include Hare (1949, 1952, 1967), and, more generally, any analysis of the imperative operator as a species of deontic necessity modal (and, potentially—depending on how the details of his idea are developed—Geach (1958)). Customarily, the proposal is motivated by appeal to *fulfillment-conditions*, which are the imperative analogue of truth- or satisfaction-conditions. Along these lines, the basic semantic relation in imperative logic—the imperative analogue of satisfaction—is conceived as fulfillment. We understand an imperative sentence $!\phi$ as being fulfilled (in a model) iff its complement indicative ϕ is true (in that model). If, and only if, a sequence of formulae $\phi_1, ..., \phi_n$ entails ψ is it the case that the fulfillment of $!\phi_1, ..., !\phi_n$ guarantees the fulfillment of $!\psi$. The final step in motivating the logic is the proposal that imperative entailment should be understood in terms of a guarantee of fulfillment: a "premise" sequence of imperatives entails a "conclusion" imperative iff the fulfillment of each premise imperative guarantees the fulfillment of the conclusion imperative. The parallels to the ordinary understanding of the entailment relation (in terms of a guarantee of truth) are obvious and will not be belabored.

The account is incomplete, along several dimensions. Most obviously, a statement of conditions on an imperative entailment relation is *not* a fully worked out semantics. In particular, we are missing an explicit account of the semantic contribution (if any) of the imperative operator to the meanings of imperative formulae. The proposal is, as it stands, compatible with treating the imperative operator as a species of deontic necessity modal or semantically null speech-act operator (among sundry other possibilities). We bracket the issue for now, but take it up again later.

Additionally, we lack any account of the proper interaction between imperatives and indicatives in argument; we have only stated necessary and sufficient conditions on when

2. We spend a good deal of time on imperative truth (§3.3) and how to interpret Boolean connectives in an imperative logic (§3.4.3). To preview: truth-conditional analyses of the imperative are not totally out of the question. But the semantic apparatuses that accompany such analyses are rich enough to handle conditional imperatives in other ways, and there are reasons for preferring them. (Imperative threats are trickier, and we will have nothing to say about them in this paper.)

an *imperative* conclusion follows from a sequence of *imperative* premises. So, for instance, we have nothing to say about whether the following sorts of arguments ought to be counted as valid or invalid.

(12) Use an axe or a saw! [$!(\phi \vee \psi)$]
You will not use an axe. [$\neg\phi$]
Then use a saw! [$!\psi$]

(13) See to it that: if you go to Harlem, you take the A-train! [$!(\phi \to \psi)$]
You are going to Harlem. [ϕ]
Then take the A-train! [$!\psi$]

Although I argue shortly for a form of skepticism about such arguments, they have some prima facie appeal. Noticing this, we might try to fill part of the gap in our proposal by strengthening it.[3]

(15) $\{!\phi_1,...,!\phi_n\} \cup \{\psi_1,...,\psi_m\} \Vdash !\pi$ iff $\{\phi_1,...,\phi_n\} \cup \{\psi_1,...,\psi_m\} \vDash \pi$

There are definite difficulties doing so. For one, endorsing the strengthened proposal means giving up any hope for an analysis of the imperative operator as a species of deontic necessity modal. On none of the standard semantical treatments of deontic logic do we have, for example, that either $O(\phi \vee \psi), \neg\phi \vDash O\psi$ or $O(\phi \to \psi), \phi \vDash O\psi$. But, while the prospect of effecting a full-scale reduction of imperative logic to deontic logic might have looked appealing, this might be thought a small concession. Other things equal, it would be preferable, for example, to avoid an analysis on which imperatives are given modalized truth-conditions.

Agreed, but it would nevertheless be surprising for the valid argument forms of imperative and deontic logic to diverge so radically. And, indeed, we find that the sorts of considerations that tell against the corresponding deontic argument forms can be leveraged against their imperative kin. The deontic prohibition on detaching $O\psi$ from premises $O(\phi \to \psi)$ and ϕ is made palatable by noting that an obligation to make a material implication $\phi \to \psi$ true might have been *best* discharged by falsifying the antecedent (and similarly for the inference of $O\psi$ from $O(\phi \vee \psi)$ and $\neg\phi$) (cf. Broome 1999). Similarly, in issuing an order of the form $!(\phi \vee \psi)$, an authority might reasonably prefer that her addressee fulfill her command by seeing to it that ϕ, and demur from endorsing the command $!\psi$ in the event that $\neg\phi$.[4] The fact that ϕ, combined with the fact that she has issued the

3. This isn't quite right. We want to rule out the validity of an argument from a sequence of non-imperative formulae to an imperative conclusion—e.g., from ϕ and $\phi \to \psi$ to $!\psi$. (15) does not. This brings up the thorny question of whether/when imperatives and indicatives might imply each other. In this vein, Hare (1952: 28) defends a proposal along these lines:

(14) a. $\{!\phi_1,...,!\phi_n\} \cup \{\psi_1,...,\psi_m\} \Vdash \pi$ iff $\{\psi_1,...,\psi_m\} \vDash \pi$
(Γ implies an indicative iff the indicatives in Γ do.)
b. $\{!\phi_1,...,!\phi_n\} \cup \{\psi_1,...,\psi_m\} \Vdash !\pi$ only if $\{!\phi_1,...,!\phi_n\} \neq \emptyset$
(Γ implies an imperative only if Γ contains at least one imperative.)

The statement of conditions on $\Vdash$ in (15), then, presupposes the correctness of (14b). Geach (1958) rejects both claims. Concerning (14a), he notes a case where $!(\phi \to \psi)$, $!\neg\psi$, and ϕ seem inconsistent, so that $!(\phi \to \psi), !\neg\psi \Vdash \neg\phi$ (*If you are loyal, rise up! But do not rise; stay on your knees! Hence: you are disloyal*). This case is puzzling, but Geach is wrong about it: an agent can utter each of these commands without thinking her addressee disloyal; indeed, she might be enjoining her addressee to be disloyal. What is more mysterious is that the command *Be disloyal!* does not seem to follow from the other two commands, although the proposal we are considering predicts that it should. Geach's argument against (14b), on the other hand, looks airtight. There are ways to finesse this issue (see, e.g., Castaneda 1958) but they are not our concern here.

4. This is *not* to say that such an authority will view ψ as impermissible. Only that she might reserve positive endorsement (of the sort expressed by an imperative $!\psi$) for those ways of fulfilling the original command that are, by her lights, both permissible *and* best. We say much more about the issues raised here in the next section.

command $!(\phi \vee \psi)$, in no way commits her to endorsing the command $!\psi$. Similar remarks could be made about the conditional case.

The question of how precisely to constrain the interaction of imperatives and indicatives in argument is vexed, and we will bracket it for now (although, as we will see, endorsing a definite semantic proposal about the imperative usually involves taking a definite position on the issue). For now, we note that there is a prima facie tension between the intuitive, deontic logic-inspired argument against the strengthened statement of conditions on $\Vdash$ and conceptualizing imperative entailment as an inclusion relation among fulfillment conditions. On the assumption that $\neg\phi$, a disjunctive imperative $!(\phi \vee \psi)$ is fulfilled if and only if ψ. To put it differently: the satisfaction of $\neg\phi$ and fulfillment of $!(\phi \vee \psi)$ (say, with respect to a valuation) guarantees the fulfillment of $!\psi$. Similarly for the conditional case. More generally, when (and only when) $\phi_1, ..., \phi_n, \psi_1, ..., \psi_m \vDash \pi$, the fulfillment of $!\phi_1, ..., !\phi_n$ and the satisfaction of $\psi_1, ..., \psi_m$ will guarantee the fulfillment of $!\pi$. So, a natural extension of the motivation for the original statement of conditions on $\Vdash$ in (11) appears to positively *recommend* the strengthened statement in (15).

There are two ways to react to this fact (as well as other facts that the logic of imperative fulfillment finds difficult to explain, which we present in the following section). One is to conclude that the *correct* imperative logic—or, at least, the correct way of conceptualizing imperative logic—is not to be found in the logic of imperative fulfillment. The other is to take this fact to motivate the development of (i) alternative conceptualizations of what imperative entailment might amount to, so that we might have some rationale for resisting the strengthened statement of conditions on $\Vdash$ in (15), and/or (ii) alternative imperative logics. I will take the latter tack. I take it that the logic of imperative fulfillment is a reasonable logic of something or other, and different imperative logics can be conceptualized, designed, and used to model different phenomena about the imperative. The logic of fulfillment simply does not appear up to the job of accounting for the sense in which the argument forms in (12) and (13) are invalid.

2.3 Logics of Content

A different approach to the logic of imperatives is what I will term the "logic of content." There is undoubtedly a sense in which a logic of imperative fulfillment is a logic of content: logics of fulfillment pay attention to *what is required* or *commanded* by an imperative sentence, and construe an imperative conclusion as following from a set Γ of formulas just in case the content of the conclusion (what the conclusion requires) follows from what the imperatives in Γ require (together with the facts that the non-imperatives in Γ describe). But I wish to identify the content of an imperative sentence with something other than what the sentence requires. What precisely imperative content amounts to will depend on the specific sort of semantic theory one endorses. But what I have in mind generally is the denotation assigned by a semantic theory to an arbitrary imperative formula $!\phi$ of $\mathcal{L}_{PI}$.

Logics of content may (and do) differ significantly from one another. Some (e.g., those that identify the denotation of $!\phi$ with that of $O\phi$, where O is a species of deontic necessity operator), will endorse the statement of conditions on $\Vdash$ given in (11). Others will not. What distinguishes logics of content as a class is a particular conceptualization of imperative logic: as normative for commanding (or, more precisely, the endorsement of commands). Ordinary propositional logic is often held to be normative, in some nebulous sense, for assertion and belief: if $\phi_1, ..., \phi_n \vDash \psi$, then asserting (or believing) that $\phi_1, ..., \phi_n$ commits an agent to endorsing an assertion (or belief) that ψ. Similarly, logics of content conceptualize imperative logic as normative for endorsement of commands: if $\{!\phi_1, ..., !\phi_n\} \cup \{\psi_1, ..., \psi_m\} \Vdash !\pi$, then endorsing each member of

$\{!\phi_1, ..., !\phi_n\} \cup \{\psi_1, ..., \psi_m\}$ commits an agent to endorsing $!\pi$.[5] Logics of content generally implement this conceptualization of imperative logic by developing novel theories of the semantic content of imperative formulae (hence their name). This is, of course, a rather crude sketch. Although we could say a good deal more by way of filling it in (and illustrating with example accounts), we will save those words now. Logics of content will be a focus of our attention in this paper. Instead, we turn to applications.

Endorsing a logic of content gives one a foothold in resisting the arguments in (12) and (13) that troubled the logic of fulfillment. Indeed, it is easy to see that the deontic logic-inspired argument against their validity *assumes* the conceptualization of imperative logic characteristic of logics of content.[6] It is, of course, only a foothold. Genuine leverage over these arguments comes only with the development of a bona fide semantics. More interestingly, logics of content are well-positioned to handle the famous "paradox" about imperative logic of Ross (1941). (Indeed, doing so seems to inspire most examples of logics of content in the literature on imperatives.)

(16) $!\phi \not\Vdash !(\phi \vee \psi)$ (*Post the letter!* $\not\Vdash$ *Post or burn the letter!*)

It is well-known that the way logics of fulfillment conceptualize imperative entailment is incompatible with the felt invalidity of the Ross inference—indeed, this point goes back to Ross himself. Since any valuation fulfilling $!\phi$ also fulfills $!(\phi \vee \psi)$, any reasonable implementation of a logic of fulfillment will have it that $!\phi \Vdash !(\phi \vee \psi)$. By contrast, conceptualizing imperative logic as a logic of content supplies a rationale for formulating a logic (and theory of imperative content) in which the Ross inference is predicted invalid: intuitively, endorsing $!\phi$ (e.g., issuing a command to post a letter) does not per se commit an agent to endorsing $!(\phi \vee \psi)$.

It is worth refining this point further. We can do so by comparing the Ross inference to the following argument.

(17) See to it that if you read the book, you come see me. [$!(\phi \to \psi)$]
Read the book! [$!\phi$]
$\therefore$ Come see me! [$!\psi$]

Castaneda (1958)—applying apparently the same sort of conception of imperative entailment we take to characterize logics of content—argues that the argument form illustrated in (17) is invalid:

> [A] teacher who [issues the premise commands in (17)] has not thereby ordered or told his student to come to see him, *regardless of the student's reading of the book* (Castaneda 1958: 43-44).

Similar reasoning might be thought to explain why logics of content should not endorse the inference of $!(\phi \vee \psi)$ from $!\phi$: if teacher commands student to post the letter, teacher has not thereby committed herself to endorsing the command to post or burn it, regardless of the student's posting it.

There are two things to note here. First, if Castaneda's argument against the argument form in (17) succeeds, it appears to rule out an entire class of proposals about the semantic content of imperatives—namely, any analysis of the imperative operator as a species of

5. Endorsement is conceived as a generic pro-attitude. Endorsement of non-imperative formulas is something like belief in them. We leave the notion of imperative formula endorsement vague for now, but take up the issue again below.

6. This isn't to say that logics of content *must* invalidate (12) and (13), only that they supply a rationale for doing so. See the end of §3.2 for some discussion.

deontic necessity modal—as possible implementations of the conceptualization of imperative entailment that is characteristic of logics of content. Consider the deontic version of the **K** axiom.

(18) $O(\phi \to \psi) \to (O\phi \to O\psi)$

K is valid in the class of all frames for Standard Deontic Logic. The following is an easily proved metatheorem of SDL.

(19) $\Gamma \cup \{\phi\} \vDash \psi$ iff $\Gamma \vDash \phi \to \psi$

Together with **K**, this implies that $O(\phi \to \psi), O\phi \vDash O\psi$. If Castaneda is right about the invalidity of (17), and we aim to analyze the imperative operator as the deontic necessity operator O of SDL, we are saddled with an obvious contradiction. (Noticing that $O\phi \to O(\phi \vee \psi)$ is valid in the class of all frames for deontic modal logic, if Castaneda-esque reasoning about the Ross Paradox is correct, we have an even more direct argument to the same conclusion.)

Second, and more interestingly, Castaneda's argument involves a tacit commitment to a definite view about the content of imperatives. The *order* ϕ *regardless of* ψ locution is most naturally interpreted as expressing that both $\phi \wedge \psi$ and $\phi \wedge \neg\psi$ are permitted (when possible), and at least one required. The reason that the conclusion imperative of (17) is held not to follow from the premise imperatives is that the *content* of the imperative, according to Castaneda, is given, roughly, by a command to come see teacher regardless of the student's reading of the book. In other words, Castaneda seems to think that the conclusion imperative expresses a requirement to come see teacher and a permission to do so with or without having read the book. Since a speaker who endorses the premise imperatives of (17) in no way commits herself to permitting her addressee not to read the book (indeed, quite the opposite), a logic of content ought to dictate that the argument is invalid. (The adaptation of Castaneda's argument to the Ross inference holds that a disjunctive imperative $!(\phi \vee \psi)$ expresses a requirement to make it the case that $\phi \vee \psi$ and a permission to do so without making it the case that ϕ (assuming this to be possible). Since a speaker who issues the command $!\phi$ in no way commits herself to endorsing such a permission, a logic of content should dictate that the Ross inference is invalid.)

The notion that imperatives bear permissions as part of their content has a significant degree of historical (see, e.g., Williams 1963; Hare 1967) and contemporary (see, e.g., Aloni 2007) appeal, and there is undoubtedly something right about it. It is impossible to consistently endorse a command without being disposed to endorse some sort of permission—*minimally*, a permission that the command be fulfilled.[7] What sort of permissions should be written into the content of imperatives (consistent with the conceptualization of imperative logic characteristic of logics of content) is, however, a different matter. I want to suggest that there are ways of integrating the notion of permission into a theory of imperative content that sanction the validity of the inference in (17) and which are consistent with the shared motivations of logics of content. (There are also, as we will see shortly, ways of integrating the notion of permission which sanction the validity of the Ross inference, but they involve commitments to questionable assumptions about the permissive content of choice-offering disjunctive imperatives.)

From the vantage of a logic of content, it is clear that inferring an order to see teacher, regardless of your reading, from the premises of (17) is fallacious. Here we are in agreement with Castaneda. But glossing the conclusion of (17) with the locution *come see me regardless of your reading of the book* involves badly misrepresenting its content: what-

7. This requirement is the imperative logic analogue of the **D** axiom of SDL: $O\phi \to \neg O \neg\phi$.

ever its permissive content, it is obvious that the conclusion is silent about whether failing to read the book is permitted. On the minimal assumption that an imperative formula $!\phi$ expresses a permission that it be fulfilled, Castaneda appears to predict that the conclusion of imperative of (17) expresses a permission that is, loosely speaking, inconsistent with the requirement of the premise imperative *Read the book!*. This is problematic ground for a logic of content to occupy. Consider the following imperative argument, which any logic of content ought to rule valid.

(20) Read the book! [$!\phi$]
Come see me! [$!\psi$]
$\therefore$ Read the book and come see me! [$!(\phi \wedge \psi)$]

If the command *Come see me!* expresses a permission to do so without having read the book, then no logic of content could sanction (in view of its conceptualization of imperative entailment) the validity of this argument: agents are never committed to endorsing a command which requires that ϕ while permitting that $\neg\phi$. And yet it is clear that any agent who endorses the premise commands in this argument is committed to endorsing the conclusion.

There is a better way of presenting a worry about (17)—one which is actually consonant with the motivation for logics of content. An agent who endorses the premise imperatives of (17) might reasonably refrain from endorsing its conclusion. In, for example, a situation where her addressee fails to read the book, an agent who endorses both premise commands might reasonably demur about *issuing* the further command *Come see me!*. This is a reasonable thing to say, but note that it is tied to a particular understanding of endorsement—one concerned with an agent's communicative dispositions: roughly, an agent endorses an imperative $!\phi$ at time t if she would have no complaint, in view of her desires and beliefs at t, about issuing $!\phi$ at t.[8]

Although this is a natural way of understanding endorsement (and presumably will characterize a reasonable subclass of imperative logics of content—e.g., that endorsed by Castaneda), I want to employ a rather different sense of endorsement. The reason is simple: the logics of content we consider in this paper appear to be exploiting this different sense of endorsement, in that they all predict the inference in (17) valid. The relevant sense of endorsement is this: an agent endorses an imperative $!\phi$ at time t if the content of $!\phi$ is a suitable expression of her desires at t. Distinguish the state of affairs required by an imperative $!\phi$ (its **command content**, i.e., the state of affairs expressed by ϕ) from the states of affairs permitted by $!\phi$ (its **permissive content**). Then, very roughly, $!\phi$ is a suitable expression of an agent's desires at t iff the agent desires its command content at t and its permissive content is compatible with what she desires at t.

This is, as it stands, so abstract as to be almost useless. We clarify with examples. Suppose it is important to an agent that you read the book, and she desires that you do so. Such an agent will not endorse a command whose content is glossed as *Come see me, regardless of your reading of the book*. Such a command expresses a permission that conflicts with her desire that you read the book, and, so, does not count as a suitable expression of her desires. An agent who utters the premise commands in (17), on the other hand, desires (or at least is committed to desiring) a future in which her addressee comes to see her (indeed, if she is rational, strictly prefers one such future—the one where her addressee reads the book—to any future in which the addressee does not come to see

8. The complaint would have to concern the content of $!\phi$, as opposed to, say, the logistics of performing the utterance. This is a good time to note that much of what's going on in this section is quite rough, but generally, I hope, precise enough to enable us to draw the relevant distinctions. We're not doing conceptual analysis here—only trying to roughly taxonomize imperative logic.

her). So the content of the conclusion imperative of (17) counts as a suitable expression of her desires, although *tokening that content in an utterance* might not.[9] Supposing an agent endorses the premise imperatives in (17), then, it follows that she is committed to endorsing the conclusion imperative. Which is to say: pairing this conception of endorsement with a logic of content means endorsing the validity of the argument form illustrated in (17).

What about the Ross inference? Here we have options, depending on how we understand the permissive content of disjunctive imperatives. If we hold that the imperative *Post or burn the letter!* semantically *expresses* a permission to burn the letter (and more generally that an imperative of the form $!(\phi \vee \psi)$ expresses a permission to bring it about that ϕ and to bring it about that ψ), the class of logics of content we are considering will rule the Ross inference invalid. This is because the content of an imperative $!\phi$ being a suitable expression of an agent's desires fails to imply that the content of an imperative $!(\phi \vee \psi)$ is as well. The latter will express a permission (to bring it about that ψ) that may fail to be compatible with the desires of the agent. So endorsing $!\phi$ does not commit an agent to endorsing $!(\phi \vee \psi)$.

Alternatively, we may hold that $!(\phi \vee \psi)$ somehow conveys, without semantically expressing, a permission to bring it about that ϕ and that ψ. If that is the case, then both the permissive and command contents of $!(\phi \vee \psi)$ are exhausted by $\phi \vee \psi$, and the content of $!(\phi \vee \psi)$ *will* count as a suitable expression of the desires of an agent who endorses $!\phi$. An agent who desires that ϕ desires, inter alia, that $\phi \vee \psi$, in which case the permissive content of $!(\phi \vee \psi)$ is compatible with her desires. It follows that the content of $!(\phi \vee \psi)$ will count as a suitable expression of her desires whenever $!\phi$ does. Endorsing $!\phi$ will commit an agent to endorsing $!(\phi \vee \psi)$.

It seems to me overwhelmingly plausible that a permission to burn the letter is part of the permissive content of *Post or burn the letter!*. So, an adequate theory of the semantic content of disjunctive imperatives should predict their free choice readings by appeal to non-pragmatic mechanisms. I will not spend a great deal of time arguing for this position in this essay, although it would be possible to.[10] I will only try to provide some basic motivation for the semantic tack.

In response to the claim of Williams (1963) that the felt permissions of disjunctive imperatives are part of their semantic content (Williams suggests treating them as presuppositions), Hare (1967) argues they are in fact conversational implicatures—apparently (although he is not explicit about the point) something rather like quantity implicatures. The reasoning by which the implicatures are derived is presumably something like this:

i. If an agent desires some π such that $\pi \vDash \phi \vee \psi$ but $\phi \vee \psi \nvDash \pi$ (and it is reasonable to expect that π may be fulfilled by her addressee), then, if she is cooperative, she will not endorse the imperative $!(\phi \vee \psi)$. (Practical analogue of the Maxim of Quantity (Grice 1989): be neither more nor less action-restrictive than required.)

ii. Suppose the agent endorses $!(\phi \vee \psi)$. Then, assuming cooperativity (and that $\phi \vee \psi \nvDash \phi$ and $\phi \vee \psi \nvDash \psi$), she desires $\phi \vee \psi$ and it is not the case that: she desires ϕ or desires ψ.

iii. If the agent desires $\neg\phi$, then her desires are inconsistent unless she desires ψ. But

9. Green (1997) draws a similar distinction, although to a different purpose.

10. There exist a number of sophisticated accounts whose sole purpose is to explain free choice effects (in particuar, free choice permissions) using exclusively non-semantic (pragmatic) mechanisms. See, e.g., Kratzer & Shimoyama (2002); Schulz (2003, 2005) for sophisticated recent examples. The predominant view in the semantics literature is that free choice permissions are entailments (see, e.g., Aloni 2007; Geurts 2005; Zimmermann 2000), and presumably that view will extend to cover free choice interpretations of imperatives with disjunctive complements.

she does not desire ψ, so she does not desire $\neg\phi$. Similarly, she does not desire $\neg\psi$. So, both ϕ and ψ are compatible with what she desires.

iv. So the agent desires $\phi \vee \psi$, but does not desire any particular way of fulfilling $!(\phi \vee \psi)$. So both bringing it about that ϕ and bringing it about that ψ are acceptable ways of fulfilling $!(\phi \vee \psi)$.

The reasoning isn't either precise or airtight, but it seems plausible enough—even, dare I say, Gricean in spirit. But there are reasons for worrying about conversational implicature accounts of free choice readings of disjunctive imperatives. Note, e.g., the badness of (21a) (and the "dual" command (21b)).

(21) a. ?Post or burn the letter. But you may not burn it!
b. ?Post or burn the letter. But do not burn it!

These sorts of constructions are nearly always marked. This is surprising if the permission to burn the letter is merely a conversational implicature of the imperative. For, as Grice notes, conversational implicatures may, in general, be felicitously cancelled. Proponents of conversational implicature accounts of free choice permissions in disjunctive imperatives owe us an explanation of why permission implicatures of disjunctive imperatives are not generally felicitously cancelable.[11]

As we see below, accommodating the sort of permissive content we are defending here need not require *explicitly building permissions into* the imperative logic and semantics (although it is far preferable, I will argue, to do so). Permissive content can, in principle, serve as a rationale for a logic that invalidates the Ross inference, but that does so without giving any sort of account of or making any explicit commitments to the permissive dimension of imperatives. (See §3.5 for further discussion in connection with a Montague-Scott-style resolution of the Ross Paradox.)

2.4 Logics of Planning

Hare (1967) attempts to characterize the motivations of those who reject the statement of conditions on $\Vdash$ given in (11), for reasons having to do with the Ross Paradox. Although Hare's project is to reveal the rejection as without basis, it is clear enough that he is presupposing the conceptualization of imperative entailment that is characteristic of a logic of fulfillment, and noting that the motivations of his envisioned target diverge from it. What is interesting about Hare's argument is that the motivations he attributes to his opponents might be used to characterize an alternative conception of imperative logic—a conception common to what I will term "logics of planning."

Hare takes it that dissatisfaction with the Ross inference stems from the (he thinks mistaken) assumption that in saying that an imperative conclusion $!\psi$ follows from a series of imperative premises $!\phi_1, ..., !\phi_n$, we say that fulfilling $!\psi$ is a necessarily satisfactory way of fulfilling the obligations issued by the premise imperatives. (According to Hare and other proponents of logics of fulfillment, of course, this gets things exactly backward.)

11. Hare (1967: 315) appears to gives a dialogue in which permissions are felicitously canceled, but the case is artificial—designed to bring out the no choice interpretation of the imperative, on which the permissive content of the imperative is exhausted by $\phi \vee \psi$. I am not claiming that no choice interpretations do not exist. Rather, I claim, with Aloni (2007), that (i) there is a *semantic* difference between no choice and free choice interpretations; (ii) disjunctive imperatives are sometiems ambiguous between free choice and no choice interpretations; (iii) no choice interpretations of $!(\phi \vee \psi)$ *are* entailed by $!\phi$, while free choice interpretations are not. No choice readings are generally dispreferred, probably for Gricean reasons, and they will not occupy a central place in our discussion. Dealing with the ambiguity will, however, require complicating the semantics and/or object language. See §3.5.

On this conception of imperative entailment, it makes sense to reject the Ross inference as invalid: since fulfilling $!(\phi \vee \psi)$ is not necessarily a satisfactory way of fulfilling $!\phi$, we should endorse a logic in which $!\phi \not\Vdash !(\phi \vee \psi)$. More generally, conditions on imperative entailment should receive the following statement:

(22) $!\phi_1, \ldots, !\phi_n \Vdash !\psi$ iff $\psi \vDash \phi_1$ and ... and $\psi \vDash \phi_n$

The primary proponent of this sort of view in the literature is Anthony Kenny (see, e.g., Kenny 1966).[12] The properties of the logic associated with the statements of conditions on $\Vdash$ in (22) are not particularly interesting (as is the case with the conditions given in (11) and (15)).[13] The logic will, of course, be *non-monotonic*—in that $\Gamma \Vdash !\phi$ will not generally imply $\Gamma \cup \{!\psi\} \Vdash !\phi$—but not so in any interesting way. Non-monotonicity is just the natural consequence of reversing the direction of the entailment relation, which is essentially all that has been done here.

More interesting, for our purposes, are the intuitions that might be used to animate this approach. Kenny's general perspective on imperative logic—one common to logics of planning as a class—is essentially agent-oriented, in that it regards imperatives as primarily encoding information that agents use to structure their practical reasoning (planning). Compare logics of content, which are essentially issuer-oriented: they treat imperatives as primarily encoding information about the desires of agents who issue them. There are a variety of ways to implement this general perspective. Kenny does so as follows.

> The logic of satisfactoriness consists of the rules which ensure that *in practical reasoning* we never pass from a fiat [i.e., plan] which is satisfactory for a particular purpose to a fiat which is unsatisfactory for that purpose (Kenny 1966: 72).

According to Kenny, a proper logic of imperatives allows drawing an imperative conclusion $!\phi$ from a set Γ of imperative premises just in case the plan associated with $!\phi$ (roughly, the plan to bring it about that ϕ) is a satisfactory *implementation* of the plans associated with Γ. Intuitively, what we have here is a logic designed to model the implementation of general, higher-order plans by way of specific, lower-order plans in a rational agent. It is clear enough why we would want such a logic to be non-monotonic: some plan Π may be a satisfactory implementation of a plan Π', although strengthening (i.e., adding requirements to) Π' might easily destroy this.

There is, however, a class of natural agent- and planning-oriented logics that (i) do not focus their attention on the implementation of plans and (ii) preserve the monotonicity of $\Vdash$. What I have in mind are conceptions of imperative logic that characterize it as in the business of expressing or generating *higher-order constraints (and freedoms) on planning activities*—or on the mathematical structures we use to model such activities—of rational agents.[14] On such conceptions, the fundamental semantic relation is something like requirement (or being in force) in view of constraints in force on an agent's planning. An imperative conclusion $!\phi$ may be drawn from a set Γ of imperative premises just in case, roughly, Γ contains at least as much *practical planning content* as $!\phi$. Just in case, that is to say, whenever the constraints on planning expressed by the premise imperatives are required or in force, the constraints on planning expressed by $!\phi$ are required or in force.

12. Geach (1966) also voices his support.
13. The matter of integrating non-imperatives into the logic remains, but we will not pursue it here.
14. In particular, I have in mind the dynamic approaches of Mastop (2005); Charlow (2008a); Veltman (2008). I will present such an account in §5 of this essay.

This is a vague sketch and all of the key notions remain undefined. But, since we develop an account along these lines in the second half of the paper, we will save precision for later. Three brief notes, however. First, how such an account handles the Ross Paradox will obviously depend on how the notion of practical content gets cashed out. An account which builds free choice permissions into practical content will make different predictions about the (in)validity of the Ross inference than one which does not. Second, I think it is clear enough that the entailment relation characterized by such a logic will be monotonic.

Third, and most interestingly, although it is clear that logics of planning and logics of content (as I have characterized them) have divergent motivations, it is not clear that a concrete example of a logic of planning would have to be distinct, in any deep sense, from a concrete example of a logic of content. Differences in motivation need not manifest as differences in the semantic analysis of imperative formulae, and the class of argument forms predicted valid by a logic of content might coincide perfectly with the class predicted valid by a logic of planning. Differences will tend to depend on two factors: how the logical formalisms of the respective classes ultimately (i) cash out the notions of content that figure in their motivations (command/permissive content for logics of content; practical content for logics of planning) and (ii) characterize the behavior of non-imperative formulas in imperative inference.

Concerning (i): the logics of planning we develop in this paper will be dynamic, in one of two senses: they focus on planning behavior in time or on the changes that updating a cognitive state with a series of formulas (imperative and otherwise) induces. Insofar as these incarnations of logics of planning direct their focus at the planning behavior of the imperative "addressee" (either the constraints that "in force" commands impose on planning behavior or the changes that imperatives have the capacity to induce in an agent's cognitive state), they may appear to be *essentially* planning-oriented. Nevertheless, I shall argue (§4.13) that there are ways of effecting a rapprochement between the conception of imperative logic as a logic of content and the conception of imperative logic as a logic of planning (although things get a bit trickier when we shift our attention to update semantics, on account of the peculiarities of imperative inference in an update logic). Command/permissive content and practical planning content can be seen as two sides of the same coin. Because the formalism we develop in the course stating a logic of planning is rather more sophisticated than the formalism we develop in the course of stating a logic of content (capable of representing ordered commands, accounting for the contrast between stable/ephemeral commands, giving a semi-realistic treatment of action and planning in time), rapprochement will appear to recommend using the more sophisticated formalism in stating a logic of content.

Concerning (ii): it is natural in certain logics of planning—those concerned with modeling cognitive update in accordance with commands—to treat the argument forms of (12), (13), and (17) as valid. Consider (12): updating a cognitive state with a command to bring it about that $\phi \rightarrow \psi$ and the information that ϕ constrains the plans of the agent. To obey the command, in view of what she knows, the agent will have to bring it about that ψ. But this is *not* an inference a logic of content is necessarily comfortable with: as argued above, endorsing $!(\phi \rightarrow \psi)$ and ϕ does not necessarily commit an agent, *qua* issuer of imperatives, to endorsing $!\psi$. While subjects are constrained to obey, authorities are not constrained to prefer obedience. Or consider (17): updating a cognitive state with $!(\phi \rightarrow \psi)$ and $!\phi$ also constrains the agent's plans. If she fails to see to it that ψ, she violates at least one of her obligations. While there are conceivable logics of content that validate all of these argument patterns (see the end of §3.2), the primary examples of such do not. Nevertheless, this is not, I shall suggest, a difference to be accounted for by appeal to the distinction between logics of content and logics of planning. Certain logics of planning—those which are dynamic in the sense of focusing on planning behavior in time, rather than on modeling cognitive update—fail to validate precisely these argument patterns. The difference is more naturally accounted for by the special properties of update-semantic treatments of the imperative. What the difference means, if anything, is that it may not be possible to construe the logic that arises from the update-semantic treatment of the imperative as anything but a logic of planning.

2.5 Conclusion

We have spent a large amount of time taxonomizing imperative logics according to their understandings of the subject matter of imperative logic and the nature of imperative entailment. We did so informally (and at times with a good deal of imprecision). Nevertheless, we were able to draw out some interesting logical properties shared by certain classes of imperative logics. There is a fairly clear menu of options for the imperative logician to choose from. The remainder of this paper is devoted to examining concrete logics that implement the motivations we have been detailing (with the exception of those that characterize logics of fulfillment). We begin by considering a group of logics of content, the primary examples of which give a semantics for the imperative operator in terms of a semantics for a deontic modal operator.

3 IMPERATIVE LOGIC AS DEONTIC LOGIC

Deontic modal analyses of the imperative—by which I mean to include both analyses that treat the imperative operator as literally a deontic modal operator, and those which regard the operators as merely having similar inferential properties—have a good deal to recommend them. As we shall see, when suitable technical complications are introduced, they can furnish analyses of conditional imperatives (Schwager 2006) and the Ross Paradox (Aloni 2007) which (i) handle the semantics in terms of a relatively familiar and well-understood model theory and (ii) appear to be consonant with the motivations of logics of content. In short, developing a deontic modal analysis of the imperative is an attractive and perspicuous way of implementing the logician of content's conception of imperative logic.

3.1 Technical Preliminaries

We begin with some familiar technical apparatus, developing it and adding more along the way as needed. This development is rigorous (in a way that existing accounts of imperatives as modals often are not) and pieces are added gradually and with care. The gain in precision is, I hope, worth the technical expense. Much of the apparatus is self-explanatory, and, in the interests of brevity, expository remarks are kept to a minimum.

3.1.1 Kripke Semantics

The Basic Deontic Propositional Language $\mathcal{L}_{DL}$ is defined in the usual way. (We take the deontic necessity operator O as primitive, so that the deontic permissibility operator $P := \neg O \neg$.) Models for $\mathcal{L}_{DL}$ are defined standardly, as follows.

(23) $\mathcal{M} = \langle W, R, V \rangle$
W is the universe (world-space)
$R \subseteq W \times W$ is an accessibility relation on W
V is an assignment of subsets of W to the atoms of $\mathcal{L}_{DL}$

The accounts we consider in this section interpret O bouletically, so that $O\phi$ reads roughly as *it must, in view of what is desired, to be that* ϕ. Intuitively, $\langle w,v\rangle \in R$ iff v is compatible with what is desired at w. The set $\{v \in W \mid wRv\}$ gives the set of ideal (in view of what is desired at w) worlds. We adopt for now the standard axiomatization (**KD**) of SDL. The semantics for Boolean formulas is classical, the semantics for modal formulas standard. Let $\mathcal{M} = \langle W,R,V\rangle$. Then:

(24) a. $\mathcal{M},w \vDash_{\mathcal{L}_{DL}} p$ iff $w \in V(p)$
b. $\mathcal{M},w \vDash_{\mathcal{L}_{DL}} \neg\phi$ iff $\mathcal{M},w \nvDash_{\mathcal{L}_{DL}} \phi$
c. $\mathcal{M},w \vDash_{\mathcal{L}_{DL}} \phi \vee \psi$ iff $\mathcal{M},w \vDash_{\mathcal{L}_{DL}} \phi$ or $\mathcal{M},w \vDash_{\mathcal{L}_{DL}} \psi$
d. $\mathcal{M},w \vDash_{\mathcal{L}_{DL}} O\phi$ iff $\forall v \in W : wRv \Rightarrow \mathcal{M},v \vDash_{\mathcal{L}_{DL}} \phi$

3.1.2 Satisfaction at Contexts

For the logician of content's purposes, it is handy to view satisfaction as relative to contexts, in addition to models and worlds. The logician of content, as I have described her, is interested in formalizing authority (speaker) commitments in view of authority (speaker) desires. Contexts are natural bearers of information about the identity and desires of their speakers. Implementing this vision requires a concomitant revision in the semantics. Models are reconceived as follows:

(25) $\mathcal{M} = \langle D,W,\mathcal{R},\mathcal{C},V\rangle$
$D = \{i_1,...,i_n\}$ is a set of individuals
$\mathcal{R} = \{R_j \subseteq W \times W \mid i_j \in D\}$ is a set of accessibility relations for $i_j \in D$
$\mathcal{C} \subseteq D \times D \times \mathcal{R}$ is a set of contexts

(26) A context $c \in \mathcal{C} = \langle s_c, a_c, R_{s_c}\rangle$
s_c is the speaker of c
a_c is the addressee of c
R_{s_c} is the accessibility relation for s_c

A logician of content, interested as she is in speaker commitments, will naturally endorse the following revision of the satisfaction conditions.

(27) a. $\mathcal{M},c,w \vDash_{\mathcal{L}_{DL}} p$ iff $w \in V(p)$
b. $\mathcal{M},c,w \vDash_{\mathcal{L}_{DL}} \neg\phi$ iff $\mathcal{M},c,w \nvDash_{\mathcal{L}_{DL}} \phi$
c. $\mathcal{M},c,w \vDash_{\mathcal{L}_{DL}} \phi \vee \psi$ iff $\mathcal{M},c,w \vDash_{\mathcal{L}_{DL}} \phi$ or $\mathcal{M},c,w \vDash_{\mathcal{L}_{DL}} \psi$
d. $\mathcal{M},c,w \vDash_{\mathcal{L}_{DL}} O\phi$ iff $\forall v \in W : wR_{s_c}v \Rightarrow \mathcal{M},c,v \vDash_{\mathcal{L}_{DL}} \phi$

The key clause is the clause for modal formulas: $O\phi$ is satisfied in a model $\mathcal{M}$ at a context-world pair $\langle c,w\rangle$ just in case ϕ is required by what the speaker of c desires at w. (Note the following convention: if $\mathcal{M} = \langle D,W,\mathcal{R},\mathcal{C},V\rangle$, then whenever we write $\mathcal{M},c,w \vDash \phi$, it is understood that $c \in \mathcal{C}$ and $w \in W$.)

3.1.3 Ordering-Source Semantics

The accessibility semantics, as stated above, is a bit too coarse for our purposes. Most obviously, the semantics presupposes that for any context-world pair $\langle c,w\rangle$, there is in the set of "admissible" worlds (i.e., W) at least one w-accessible world v, i.e., at least one world that is ideal in view of the desires at w of the speaker of c. (Indeed, this is precisely what the **KD** axiomatization of the logic guarantees.) This idealization is harmless in giving a semantics for the basic deontic language. It is a special case of the so-called "Limit Assumption" (Lewis 1973), but I will be happy to work within the Limit Assumption in the confines of this paper (see fn17). Nevertheless, complicating our languages (and the semantics for those languages) to handle conditional imperatives will motivate abandoning it. See §3.4.4 for further discussion.[15]

It goes without saying that the ordering-source semantics formulated here is in all of its essentials drawn from the classic semantics for modals developed in Kratzer (1981). We reconceive models by eliminating accessibility relations and complicating the context parameter:

(28) $\mathcal{M} = \langle D, W, \mathcal{C}, V \rangle$
$\mathcal{C} \subseteq D \times D \times \mathcal{B} \times \mathcal{B}$ is a set of contexts
$\mathcal{B} = \{b \mid b : W \mapsto 2^{2^W}\}$ is a set of conversational backgrounds

(29) A context $c \in \mathcal{C} = \langle s_c, a_c, f_c, g_c \rangle$
$f_c \in \mathcal{B}$ is the modal base in c
$g_c \in \mathcal{B}$ is the ordering-source in c

Modal bases and ordering sources are typed as functions from worlds to sets of propositions. Together they define the domain over which the deontic modal O quantifies. Taking the intersection of the set of propositions in the modal base will characterize a set of admissible worlds (admissible in view of the information relevant at a context). Ordering sources, on the other hand, are used to rank admissible worlds in view of the desires of the speaker at a context: $g_c(w)$ yields the set of propositions that the speaker of c desires to be true at w. The modal's domain of quantification is the set of *best* (rather than ideal) admissible worlds, relative to the ordering source.

Formally, we use the ordering-source to define a preorder on W.

(30) $v \leq_{g_c(w)} u$ iff $\{p \in g_c(w) \mid u \in p\} \subseteq \{p \in g_c(w) \mid v \in p\}$

The domain of the operator O at a context-world pair $\langle c, w \rangle$, $\mathit{min}(f_c(w), \leq_{g_c(w)})$,[16] is defined as the set of admissible worlds v such that no admissible world u is strictly better (in view of what the speaker of c desires at w) than v.

(31) $\mathit{min}(f_c(w), \leq_{g_c(w)}) =$
$\{v \in \bigcap f_c(w) \mid \forall u \in \bigcap f_c(w) : u \leq_{g_c(w)} v \Rightarrow v \leq_{g_c(w)} u\}$

Two noteworthy consequences of this definition.

(32) a. **Realism.**
$\mathit{min}(f_c(w), \leq_{g_c(w)}) \subseteq \bigcap f_c(w)$
b. **Monotonicity.**
For any $\mathcal{P}$, $\mathcal{P}' \subseteq 2^W$ such that $\bigcap \mathcal{P} \subseteq \bigcap \mathcal{P}'$:
if $v \in \bigcap \mathcal{P}$ and $v \in \mathit{min}(\mathcal{P}', \leq_{g_c(w)})$, then $v \in \mathit{min}(\mathcal{P}, \leq_{g_c(w)})$

Realism means the selection of best worlds (according to the ordering source) never takes us beyond admissible worlds. Monotonicity means that if a world is best (according to an ordering source) in a set of admissible worlds, it remains best in a smaller set (according to the same ordering). Reducing competition cannot worsen a world's position in the ranking. This follows from the monotonicity of the partial identity function $f : \bigcap \mathcal{P}' \mapsto \bigcap \mathcal{P}$ with respect to $\leq_{g_c(w)}$.

Satisfaction conditions for Boolean formulas are given as before. The satisfaction

15. We could give an ordering-source semantics with accessibility relations, but it's cleaner if we do not.
16. This handy notation is cribbed from Gillies (2007).

conditions for modal formulas receive the following statement:[17]

(35) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DL}} O\phi$ iff $\forall v \in \mathit{min}(f_c(w), \leq_{g_c(w)}) : \mathcal{M}, c, v \vDash_{\mathcal{L}_{DL}} \phi$

On the ordering-source semantics, $O\phi$ is satisfied in a model $\mathcal{M}$ at a context-world pair $\langle c, w\rangle$ just in case ϕ is satisfied at all the best (in view of what the speaker of c desires at w) admissible worlds.

Before moving on, a final piece of handy notation: the introduction of an interpretation function $[\![\cdot]\!]^{\mathcal{L}}$ yielding the set of worlds satisfying a formula of language $\mathcal{L}$ in a model at a context: if $\mathcal{M} = \langle D, W, \mathcal{C}, V\rangle$, then $[\![\phi]\!]^{\mathcal{L}}_{\mathcal{M},c} = \{w \in W \mid \mathcal{M}, c, w \vDash_{\mathcal{L}} \phi\}$. The language superscript will generally be omitted.

3.2 Semantics for $\mathcal{L}_{PI}$

Giving the semantics for the Basic Propositional Imperative Language in terms of the semantics for the Basic Deontic Propositional Language is, for the most part, trivial: we give satisfaction conditions for formulae of $\mathcal{L}_{PI}$ in terms of satisfaction conditions for formulae of $\mathcal{L}_{DL}$. (We bracket conditional imperatives for now.) We interpret $\mathcal{L}_{PI}$ using an ordering-source semantics and identify the class of models for $\mathcal{L}_{DL}$ with the class of models for $\mathcal{L}_{PI}$. Satisfaction conditions are formulated as follows.

(36) For all $\phi \in \mathcal{L}_{PI}$:
a. If ϕ is Boolean, $\mathcal{M}, c, w \Vdash_{\mathcal{L}_{PI}} \phi$ iff $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DL}} \phi$.
b. If $\phi = !\psi$ (for $\psi \in \mathcal{L}_P$), $\mathcal{M}, c, w \Vdash_{\mathcal{L}_{PI}} \phi$ iff $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DL}} O\psi$

Imperative validity is defined in terms of satisfaction-preservation in all models and at all context-world pairs. Let $\phi_1, ..., \phi_n, \psi$ be arbitrary formulae of $\mathcal{L}_{PI}$. Then:

(37) $\phi_1, ..., \phi_n \Vdash_{\mathcal{L}_{PI}} \psi$ iff $\forall \mathcal{M}cw$:
$\mathcal{M}, c, w \Vdash_{\mathcal{L}_{PI}} \phi_1 \wedge ... \wedge \phi_n \Rightarrow \mathcal{M}, c, w \Vdash_{\mathcal{L}_{PI}} \psi$

It is obvious that the logic will endorse the statement of conditions on $\Vdash$ given in (11), although not the statement of (15), which is what we would expect (so far, anyway) from a logic of content.

For certain purposes, however, we might *want* to formulate a logic of content satisfying (15). One way of conceiving of imperative necessity consistent with a logic of content is necessity in view of what the speaker desires *and knows to be true* (as opposed to what she merely desires). Such a logic could be rather easily developed by *dynamicizing* the entailment relation, so that non-imperative formulae on the left of $\Vdash$ are "added" to the modal base against which the satisfaction of imperative formulae is checked. Taking this tack involves representing non-imperative formulae on the left of $\Vdash$ hypothetically as epistemic necessities—as holding at all worlds compatible with the relevant information.

(38) Definition. Fix $\mathcal{M} = \langle D, W, \mathcal{C}, V\rangle$, and let $c = \langle s_c, a_c, f_c, g_c\rangle$. Then:
$c + \Gamma =_{df} \langle s_c, a_c, f'_c, g_c\rangle$, where $f'_c(w) = f_c(w) \cup \{[\![\phi]\!]_{\mathcal{M},c}\}$, for all $w \in W$, $\phi \in \Gamma$.

17. This is a simplified presentation. The official statement of Kratzer (1981) is this:

(33) $I_{\leq_{g(w)}}(f_c(w), v) =_{df} \{z \in \bigcap f_c(w) : z \leq_{g(w)} v\}$

(34) $\mathcal{M}, c, w \vDash O\phi$ iff $\forall u \in \bigcap f_c(w) : \exists v \in I_{\leq_{g(w)}}(f_c(w), w)) : \forall t \in I_{\leq_{g(w)}}(f_c(w), v)) : \mathcal{M}, c, t \vDash \phi$

The official view avoids the Limit Assumption—that $\mathit{min}(f_c(w), \leq_{g_c(w)}) \neq \emptyset$. See Swanson (2008b) for discussion. For simplicity, we assume that W is finite, so the Limit Assumption does no harm. See Kolodny & MacFarlane (2008: 18) for a defense of the Limit Assumption in the deontic context.

(39) Let $\Gamma^! = \{!\phi_1, ..., !\phi_n\}$ be a set of imperative formulae of $\mathcal{L}_{PI}$,
$\Gamma^* = \{\psi_1, ..., \psi_m\}$ be a set of non-imperative formulae of $\mathcal{L}_{PI}$,
π be an arbitrary formula of $\mathcal{L}_{PI}$.
Then $\Gamma^! \cup \Gamma^* \Vdash_{\mathcal{L}_{PI}} \pi$ iff $\forall \mathcal{M} c w$:
$\mathcal{M}, c + \Gamma^*, w \Vdash_{\mathcal{L}_{PI}} \Gamma^! \cup \Gamma^* \;\Rightarrow\; \mathcal{M}, c + \Gamma^*, w \Vdash_{\mathcal{L}_{PI}} \pi$

The entailment relation can be glossed as follows: $\Gamma^!$ and Γ^* entail $!\phi$ iff ϕ is required by the speaker's desires everywhere that Γ^* is satisfied (cf. the notion of "quasi-validity" defined in Kolodny & MacFarlane (2008: 25)). We won't investigate the properties of such a logic, although it is easy to check that it entails (15). It is a virtue of the modal semantics we have given that it is flexible enough to handle these divergent ways of developing a logic of content.[18]

3.3 *Interlude: The Incredulous Stare*

Deontic modal analyses of the imperative appear to give imperatives literal truth-conditions—an imperative formula $!\phi$ can be either *true* or *false*, depending on whether ϕ is in fact obligatory, in view of some obligation-determining object. Specifically, the family of analyses we are considering apparently has it that an imperative formula $!\phi$ of $\mathcal{L}_{PI}$ serves to state a claim about the desires of its speakers: roughly that ϕ must, in view of the speaker's desires, to be the case.

Writers too numerous to cite regard any truth-conditional semantics for imperative formulae as a non-starter. Arguments for this stance are not as common as one would hope. It is often treated as a basic, Moorean fact that imperative formulae just could not have truth-conditions. No matter how much a truth-conditonal logic for imperatives did for us, there would be a strong presumption against it. This is a dogmatic way of presenting what is a prima facie legitimate suspicion, and we can do better.

For example: suppose imperative formulae (and their analogues in natural language) have the same sort of semantic content as run-of-the-mill indicatives. The fact that indicatives express propositions allows us to do a lot of neat stuff with them. For example:

- Embed them freely under the scope of truth-functional operators.
- Embed them freely under the scope of intensional operators.
- Assert the propositions they express.
- Target the propositions they express with linguistic assent and denial.

If imperatives are also in the business of expressing propositions, the argument goes, we might expect to be able to do these same things with them. Clearly, though, we cannot.

(40) *It's not the case that: please come see me right away.

(41) *He knows that please come see me right away.

(42) *I affirm that please come see me right away.

(43) a. Please come see me right away!
b. ?That's true. I really must do that, given what you want.
c. ?That's false. I really mustn't do that, given what you want.

18. While we will not bother dynamicizing the entailment relation that is associated with later revisions of the semantics for imperative formulas, the basic idea employed here—having the domain of the modal be progressively restricted by non-imperative formulae to the left of the turnstile—can be easily extended to those cases.

What these examples seem to show is that imperatives are not interpretable (by human language processors) as expressing propositions.[19] The opponent of the truth-conditional analysis will continue that if they did express propositions, then they would be so interpretable. Conclusion: imperatives do not express propositions.

Grant that imperatives (or, more precisely, utterances thereof) are not interpretable as expressing propositions. (For short, we will say that imperatives have essentially *performative* interpretations. What precisely we mean by this will be clarified below, but roughly this means that imperatives are capable of receiving only obligation-imposing, rather than obligation-describing, interpretations.) There is yet room for resistance.

To illustrate: Schwager (2006), which endorses a speaker-relative bouletic modal semantics for the imperative operator, contends that *lexical presuppositions* of the imperative operator explain why natural language imperatives have exclusively performative interpretations. The crucial presupposition, by her lights, is that the speaker of an imperative at c affirms g_c as a good source of rules for acting.[20] This unfortuantely raises more questions than it answers. What, for example, does this sort of affirmation involve? It cannot, of course, be the sort of affirmation (i.e., of the truth of a proposition) involved in assertion. Presumably Schwager intends affirmation as its own type of speech-act—a speech-act of the sort that is incompatible with an assertoric interpretation of an utterance used to perform it. Perhaps there is something to this. But ultimately the proposal fails to explain the explanandum: imperatives are claimed to have exclusively performative interpretations *because* they can only be tokened in a performative sort of speech-act (affirmation).

While Schwager's proposal is unpersuasive, it is suggestive. Whatever the right account of the performative force of imperatives—in terms of lexical presupposition or some other pragmatic mechanism (e.g., the To-Do Lists of Isaacs & Potts (2003); Potts (2003); Portner (2004, 2008), which we tackle at length below)—it seems likely that it will be, at a minimum, *compatible* with the semantics of imperatives we have on offer. Performative force is a pragmatic notion and it is not unreasonable to think that it will have a purely pragmatic explanation—one orthogonal to a semantic account of imperative validity. Whatever the right account of performative force, we can expect to be able to graft it onto a reasonable semantics for imperatives without inconsistency, thereby combining our semantic cake-having and pragmatic cake-eating.

In this vein, it seems clear that Portner (2008) is misguided in his objections to analyses which treat the imperative operator as a modal.

> A modal which only had a performative use might as well not be called a modal at all. The performative aspect of its meaning, modeled as the addition of its prejacent to the To-Do List or in some other way, would explain everything that needs to be explained about its meaning. (Portner 2008: 363).

Portner illustrates the point by appeal to the cases of Ninan (2005) chronicling performative uses of root modals with exclusivity presuppositions (*must*, *have to*, ...). In such

19. This is actually complicated. It may be that the asterisked constructions subcategorize for complement IPs with non-null specifiers (i.e., overt subjects). If that were right, the badness of these constructions could be explained non-semantically. I do not know enough about the syntactic issues to pursue them with any seriousness here.

20. Schwager (2006) makes hay over another putative presupposition borne by the imperative operator: that the speaker is an epistemic authority about f_c and g_c (in a very technical sense not relevant for our purposes here). This presupposition is designed to induce the further presupposition that the speaker "cannot be mistaken" about whether or not the modal sentence giving the content of an imperative utterance is true—equivalently, that the speaker is an authority on what her desires require (with respect to the modal base). Schwager does not say what this has to do with performative force. There are many subject matters about which I cannot be mistaken and yet still felicitously state facts. See §5.2 for further discussion of Schwager's proposal.

cases, it seems clear that the truth-conditional semantics of the modals does not factor in the correct account of the performative force of the utterances.[21] So the truth-conditional semantics of modals is, in the relevant sense, pragmatically inert. Since, according to Portner, there is nothing to explain about the meaning of imperatives *beyond* their performative aspect (in contrast with root modals), the truth-conditional semantics does no theoretical work and ought to be abandoned. More generally, Portner is objecting to any account of imperative meaning—truth-conditional (or not), modal (or not)—that is not a theory of (or built on top of a theory of) imperative aspect.

This strikes me as shortsighted. Quite clearly, there is a good deal to explain about imperative meaning beyond performative aspect. In particular, we need a *semantic* account of the imperative—a characterization of the fundamental semantic relation for an imperative language and an account of imperative validity. If construing imperative operators in terms of deontic modality helps us with these tasks, from the vantage of one of the conceptions of imperative entailment we have described, then there is good, if defeasible, theoretical reason for viewing imperatives as having modalized truth-conditions.

It is true that on some conceptions of imperative entailment—in particular, the sort characteristic of logics of planning—it seems possible to use one and the same formal apparatus to account for performative aspect and to characterize a satisfactory imperative entailment relation, without giving imperatives modalized truth-conditions. (We develop such an account at length in §5.) In which case, there might be no compelling reason for endorsing a modal account: all the semantic heavy-lifting is done by the pragmatic apparatus. Things are not so simple as that, however. For one, it isn't obvious that a sufficiently flexible logic of content *could* be given in terms of this apparatus. To preview, the state of the art in accounting for imperatives' performative aspect has it that imperatives add to their addressees' plans or commitments (the above-mentioned To-Do List approach). Performative force is held to consist in constraining the plans of an addressee, and the accompanying logic seems to be *about* modeling higher-order constraints on planning behaviors of the addressee, rather than speaker commitments to endorse, in view of their desires. Even if we can give a logic-of-content rationale to the logic—and I will argue that we can (§4.13)—it may turn out that treating imperative logic in terms of modal deontic logic is the best way of implementing a logic of content. We argue the point in §5.5.

There is a way to mitigate the discomfort of a modal account: reconceive the meaning of the imperative "satisfaction" relation $\Vdash_{\mathcal{L}_{PI}}$ for imperative formulae, in terms of some less truthy notion. In this vein, Lemmon (1965) suggests the notion of an imperative's *being in force*, while Segerberg (1990) suggests *requirement*.[22] Supposing we take this tack, the semantics given in (36) and validity definition given in (37) will remain the same, but carry a slightly altered meaning. When $\phi \in \mathcal{L}_P$, a model $\mathcal{M}$ and context-world pair $\langle c, w\rangle$ will require $!\phi$ ($!\phi$ is in force in $\mathcal{M}$ at $\langle c, w\rangle$) just in case they satisfy the corresponding statement of obligation $O\phi$. In brief: *the imperative's requirement conditions are identified with its corresponding obligation-statement's truth-conditions*. To say $\Gamma^! \cup \Gamma^* \Vdash \phi$ (where $\Gamma^!$ is a set of imperative formulae and Γ^* a set of non-imperative formulae of $\mathcal{L}_{PI}$) is to say that whenever a model $\mathcal{M}$ and context-world pair $\langle c, w\rangle$ require each member of $\Gamma^!$ and satisfy each member of Γ^*, $\mathcal{M}$ and $\langle c, w\rangle$ require ϕ (if ϕ is imperative) and satisfy ϕ (if ϕ is non-imperative). This strategy preserves close structural parallels between the imperative logic and its deontic cousin. In particular, for any argument form (in)validated by the former, its deontic cousin—the result of uniformly substituting deontic O for imperative ! in the argument—will be (in)validated by the latter. It also seems to tactfully sidestep the

21. This is actually much too quick. See §4.13 for discussion.

22. Lemmon's suggestion is superior if we elect to accommodate permissions (which we have been ignoring). It makes good sense to say that a permission is in force, considerably less sense to say that a permission is required.

issue about imperative truth. What we have, then, is a semantics for imperatives in terms of the model theory for deontic modal logic, but which regards imperatives as expressing requirements that are either in force or not, deontic formulas as expressing descriptions of requirements that are true or not, depending on whether those requirements are in force or not.

What it does *not* sidestep is a certain sort of hazy methodological worry, one which someone like Portner would probably endorse: given the essentially performative nature of imperatives, we have, other things being equal, reason to prefer a unified account of the semantics and pragmatics of imperatives, and to disprefer semantic accounts which bear no obvious connection to the performative aspect of imperatives. I share the worry (even in its hazy state), and try to address it tentatively below. Interlude over. Now back to brass tacks.

3.4 Conditional Imperatives

Conditional (qualified) imperative constructions like (44a), which we represent schematically as in (44b), are ubiquitous in natural language.

(44) a. If your boss comes, offer her a seat!
b. $(\textit{if}\ \phi)(\textit{stit}\ \psi)$

We should like to have some way of representing conditional imperatives in an imperative logic. The menu of options for doing so is as follows.

3.4.1 Wide-Scoping

A wide-scope treatment of the conditional imperative $(\textit{if}\ \phi)(\textit{stit}\ \psi)$ is any that represents it with the $\mathcal{L}_{PI}$ logical form $!(\phi \rightarrow \psi)$. Wide-scope treatments will fail to validate argument forms like the following.

(45) If your boss comes, offer her a seat! $[(\textit{if}\ \phi)(\textit{stit}\ \psi)]$
Your boss is coming. $[\phi]$
So, offer her a seat! $[\textit{stit}\ \psi]$

It is, for now, an open question whether a logic content should validate such argument forms (although we try to close it below). More generally, it is an open question whether modus ponens for conditional imperatives should be validated in a logic of content—whether a logic of content should have it that *stit* ψ follows from ϕ and $(\textit{if}\ \phi)(\textit{stit}\ \psi)$. Formally, a modus ponens rule for conditional imperatives amounts to the following semantic constraint: if (the logical forms of) ϕ and $(\textit{if}\ \phi)(\textit{stit}\ \psi)$ are required at w, then (the logical form of) *stit* ψ is also required at w.

But no matter whether we endorse modus ponens for conditional imperatives, the wide-scoper is in trouble. Suppose we endorse it. There are two ways for the wide-scoper to accommodate this, consistent with the imperative logic we are considering. First, she might axiomatize the logic or constrain the model theory so that $O(\phi \rightarrow \psi), \phi \vDash_{\mathcal{L}_{DL}} O\psi$. But we have this only on the assumption that, for any $\mathcal{M} = \langle D, W, \mathcal{C}, V\rangle$, $w \in W$, $c \in \mathcal{C}$: $v \in \textit{min}(f_c(w), \leq_{g_c(w)}) \Rightarrow v = w$. Adopting this assumption would be disastrous—for starters, $\phi \rightarrow O\phi$ would be valid in the class of all models for the logic. Obviously this is a non-starter.

Second, she might utilize the dynamicized entailment relation ("quasi-validity") detailed at the end of §3.2. This is also a non-starter. Adopting the dynamicized entailment relation would mean *always* characterizing the inference of $!\psi$ from $!(\phi \rightarrow \psi)$ and ϕ as valid. But this would be too extreme. *Qua* logicians of content, and assuming the validity of modus ponens for conditional imperatives, we would like to have *available* a logic which represents inferences of the form $(if\ \phi)(stit\ \psi), \phi / stit\ \psi$ as valid, without having it that $!\psi$ should *always* follow from $!(\phi \rightarrow \psi)$ and ϕ. Recall cases (12) and (13). If we wide-scope, it seems there is no such logic to be had.

3.4.2 Interlude: Modus Ponens

The wide-scoper can object: as Gillies (2008); Kolodny & MacFarlane (2008) show, modus ponens is not in general valid for natural language conditionals. So we should not expect its imperative analogue to be valid, in which case the wide-scoper might think she is off the hook. Consider the following case.

> Ten miners are trapped either in shaft A or in shaft B, but we don't know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed (Kolodny & MacFarlane 2008: 1).

In such a case, it is reasonable for an authority to endorse any of the imperatives in (46) (indeed, all of them simultaneously), while failing to endorse any of the imperatives in (47).[23]

(46) a. Block neither shaft. $[stit(\neg bl\_A \wedge \neg bl\_B)]$
b. If they're in A, block A. $[(if\ in\_A)(stit\ bl\_A)]$
c. If they're in B, block B. $[(if\ in\_B)(stit\ bl\_B)]$

(47) a. Block A. $[stit\ bl\_A]$
b. Block B. $[stit\ bl\_B]$
c. Block one of the shafts. $[stit(bl\_A \vee bl\_B)]$

The lesson of the case is that a logic of content's analysis of conditional imperatives should not have it that any of the following semantic relationships hold.

- When ϕ_1 and ϕ_2 partition the set of worlds compatible with the modal base at any world of evaluation, $(if\ \phi_1)(stit\ \psi_1)$ and $(if\ \phi_2)(stit\ \psi_2)$ entail $stit\ \psi_1$ or $stit\ \psi_2$ or $stit(\psi_1 \vee \psi_2)$
- $(if\ \phi_1)(stit\ \psi_1)$, $(if\ \phi_2)(stit\ \psi_2)$, $\phi_1 \vee \phi_2$ entail either $stit\ \psi_1$ or $stit\ \psi_2$ or $stit(\psi_1 \vee \psi_2)$

Violating the first of these relationships is in direct tension with intuitions about the miners case. Violating the second would mean that adding the premise $in\_A \vee in\_B$ to premises $(if\ in\_A)(stit\ bl\_A)$ and $(if\ in\_B)(stit\ bl\_B)$ would allow inferring one of the imperatives in (47). In which case, on a logic of content, endorsing $in\_A \vee in\_B$, $(if\ in\_A)(stit\ bl\_A)$, and $(if\ in\_B)(stit\ bl\_B)$ would commit an agent to endorsing one of the imperatives in (47). And this also violates intuitions about the miners case.

23. Interpreting the data in this case is actually a complicated endeavor. I assume one plausible interpretation here. Kolodny and Macfarlane do a good job disposing of resistance to similar intuitions about sentences identical save for the uniform replacement of the imperative *stit* with the deontic *ought*. I won't recapitulate here, but it's clear that their arguments can be applied mutatis mutandis in defense of the intuitions I am insisting on here.

Alas, modus ponens for conditional imperatives predicts the second of these relationships to hold. Suppose $(\textit{if}\ \phi_1)(\textit{stit}\ \psi_1)$ and $(\textit{if}\ \phi_2)(\textit{stit}\ \psi_2)$ are required and $\phi_1 \vee \phi_2$ satisfied at w. Then w satisfies either ϕ_1 or ϕ_2. In either case modus ponens for conditional imperatives has it that one of *stit* ψ_1 or *stit* ψ_2 is required at w. So it cannot be right.

A fair point—one we return to below—but it does not help the wide-scoper. Suppose ϕ_1 and ϕ_2 partition the set of worlds compatible with the modal base at any world of evaluation (we may suppose this is what is going on in Kolodny and MacFarlane's case) and $O(\phi_1 \rightarrow \psi_1)$ and $O(\phi_2 \rightarrow \psi_2)$ are satisfied. It follows that $O(\psi_1 \vee \psi_2)$ is satisfied. Why? When all the best worlds satisfy $\phi_1 \rightarrow \psi_1$, $\phi_2 \rightarrow \psi_2$, and $\phi_1 \vee \phi_2$, all the best worlds satisfy $\psi_1 \vee \psi_2$. In which case $O(\psi_1 \vee \psi_2)$ is satisfied and, by the imperative clause of (36), $!(\psi_1 \vee \psi_2)$ is required. A logic of content's analysis of conditional imperatives should not have it that, when ϕ_1 and ϕ_2 partition the set of worlds compatible with the modal base at any world of evaluation, $(\textit{if}\ \phi_1)(\textit{stit}\ \psi_1)$, $(\textit{if}\ \phi_2)(\textit{stit}\ \psi_2)$, and $\phi_1 \vee \phi_2$ entail $\textit{stit}(\psi_1 \vee \psi_2)$.[24]

Here is the state of play. It seems likely that modus ponens for conditional imperatives is not generally valid—that we cannot *always* have it that $(\textit{if}\ \phi)(\textit{stit}\ \psi), \phi$ entail *stit* ψ. (Whether or not this is right, the wide-scoper is out of luck.) It may yet be possible to characterize the inference in (45) as valid, in some sense or other, without presupposing the general validity of modus ponens for conditional imperatives. We return to the issue in §3.4.4.

3.4.3 Narrow-Scoping

A narrow-scope treatment of the conditional imperative $(\textit{if}\ \phi)(\textit{stit}\ \psi)$ is any that represents it with the logical form $\phi \rightarrow !\psi$. Recall that such formulae are not well-formed formulae of the Basic Propositional Imperative Language. So adopting a narrow-scope analysis of the conditional imperative will require redefining the object language. There are several options for doing so.

Implicitly, we have been taking $\rightarrow$ to be defined in terms of $\neg$ and $\vee$: $\phi \rightarrow \psi := \neg\phi \vee \psi$. If, in redefining the imperative language, we elect to keep things this way, the logical form of the conditional imperative $(\textit{if}\ \phi)(\textit{stit}\ \psi)$ will abbreviate the formula $\neg\phi \vee !\psi$. Earlier, we provisionally objected to this sort of design, on the grounds that it would be committed to the ambiguous interpretation of Boolean connectives. But this point needs to be refined, since understanding the imperative analogue of satisfaction in terms of requirement involves *assuming the ambiguity into the metalanguage* (in particular, into the meaning of $\Vdash$). The semantic interpretation of $\vee$ is constant and unambiguous: for all ϕ, $\psi \in \mathcal{L}_{PI}$, $\mathcal{M}, c, w \Vdash \phi \vee \psi$ iff $\mathcal{M}, c, w \Vdash \phi$ or $\mathcal{M}, c, w \Vdash \psi$. What this metalinguistic condition actually amounts to will depend on the syntactic types (imperative or non-imperative) of ϕ and ψ, but it is not clear why that would be worrisome.

But there is a related worry. The semantics may interpret $\vee$ uniformly, but *interpreters* do not. When ϕ and ψ are non-imperative, it is natural to interpret $\phi \vee \psi$ in terms of its metalinguistic (English) satisfaction conditions: what $\phi \vee \psi$ expresses is that ϕ and/or that ψ. In contrast, suppose that $\psi = !\pi$. We cannot interpret $\phi \vee \psi$ in terms of its metalinguistic requirement conditions. There is no accessible reading of $\phi \vee !\pi$ on which it expresses that ϕ and/or $!\pi$; indeed, there is no accessible reading of this disjunction at all.[25]

24. Can an account of the permissive content of $!(\psi_1 \vee \psi_2)$ save the wide-scoper? No. The problem stems from the command, not the permissive, content of the imperative. It commands something that an authority may often reasonably demur from commanding.

25. Due to the fact that performatives are not embeddable under the scope of sentential connectives in any natural language (with the possible exception of negation in languages like Dutch; see Veltman (2008)). Portner (2008: 379-80) cites examples of embeddings of imperatives under verbs which take a sentential complement in Korean. While interesting, this is not germane to the issue here.

I assume it as condition of adequacy on the imperative object language that it not contain formulas which human language processors are unable to interpret according to their intended interpretations (and especially that object language representations of natural language imperative constructions not be uninterpretable in this way). Other authors (e.g., Segerberg 1990) do not make this assumption, allow imperative formulae to embed relatively freely, and are happy with treating hypothetical imperatives in terms of Boolean $\rightarrow$. Their primary interest, however, is not in designing a language and model theory for the eventual purpose of doing serious, formal analysis of natural language. Insofar as ours is, it is reasonable to restrict the embeddability of imperative formulae in the object language in this way.

For a narrow-scoper, this means designing an imperative language and semantics with the following properties: (1) $\rightarrow$ is not defined in terms of other Boolean connectives, nor other connectives in terms of it; (2) imperative formulae cannot occur in antecedent position[26]; (3) $\phi \rightarrow !\psi$ is assigned a non-Boolean interpretation, since there is no accessible reading of $\phi \rightarrow !\psi$ on which it expresses that $\neg\phi$ and/or $!\pi$. This is a daunting project, and will apparently involve either failing to treat $\rightarrow$ univocally (or else failing to treat it as a Boolean connective altogether). There are dynamic accounts that implement something like the narrow-scoper's approach (Asher & Lascarides 2003; Mastop 2005; Potts 2003), according to which $\rightarrow$ receives a uniform update potential—one which constrains information states so that worlds satisfying the antecedent and compatible with is known meet a certain condition—satisfying or requiring the consequent, as the case may be. But I have no idea whether the project could be carried out in a reasonable way in the sort of static, speaker-focused, model-theoretic framework we are currently working within.[27]

3.4.4 Two-Place Imperative Operators

Wide-scoping and narrow-scoping have empirical and conceptual drawbacks, respectively. A different approach—one which analyzes imperatives in terms of the standard account of modals in natural language as two-place generalized quantifiers—can do better. Our concern is, as we have said, primarily with formal languages in this paper. Taking this cue from the semantics of natural language turns out, however, to have tangible benefits.

Lewis (1975) and Kratzer (1991) present data suggesting that the general function of *if*-clauses in natural language conditionals is to restrict the domain of a two-place generalized quantifier occuring (sometimes overtly, sometimes covertly) in *then*-clauses. The surface syntax of natural language conditionals is misleading:

> The history of the conditional is the history of a syntactic mistake. There is no two-place *if ... then* connective in the logical forms for natural languages. *If*-clauses are devices for restricting the domains of various operators. Whenever there is no explicit operator, we have to posit one Kratzer (1991: 656).

In order to make the deontic logic we're working with reflect this insight, we need to reconstrue O as two-place—its first argument restricts the set of accessible worlds, while its second gives the condition that must be satisfied throughout the domain.

26. Possible complicating data: apparent *only if* imperatives, like *Shoot only if ordered*. There is a real question about whether such constructions are really imperative; informants tend to hear *Shoot only if ordered* as *you may shoot only if you are ordered*. Insofar as the sentence is interpretable as an imperative, though, we can handle the sentence without sanctioning imperative antecedents: we apparently lose nothing by treating *Shoot only if ordered* in the same way we treat *If not ordered, do not shoot*.

27. For a general critique of these sorts of approaches in the dynamic context, see Charlow (2008b).

This change is most naturally effected syntactically, i.e., via a change in the language. We *replace* the modal clause in the recursive definition of $\mathcal{L}_{DL}$ with the following clause. The resulting language is termed $\mathcal{L}_{DLK}$.

(48) If $\phi, \psi \in \mathcal{L}_{DLK}$, then $O(\phi)(\psi) \in \mathcal{L}_{DLK}$

A two-place permission operator P may be defined in terms of O, as follows: $P(\cdot)(\cdot) := \neg O(\cdot)(\neg\cdot)$.

The idea is that a deontic natural language conditional with surface form $(\textit{if } \phi)(\textit{must } \psi)$ will generally be associated with the $\mathcal{L}_{DLK}$ logical form $O(\phi)(\psi)$. When no restrictor is explicit—when we have a bare natural language modal of surface form *must* ϕ—the sentence is associated with the $\mathcal{L}_{DLK}$ logical form in which the domain of O is vacuously restricted: $O(\top)(\phi)$, where $\top$ is a classical propositional tautology.[28] The satisfaction conditions for modal formulas are restated as follows.

(49) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLK}} O(\phi)(\psi)$ iff
$\forall v \in \textit{min}(f_c(w) \cup [\![\phi]\!]_{\mathcal{M},c}, \leq_{g_c(w)}) : \mathcal{M}, c, v \vDash_{\mathcal{L}_{DLK}} \phi$

$O(\phi)(\psi)$ is satisfied in a model $\mathcal{M}$ at a context-world pair $\langle c, w\rangle$ just in case ψ is satisfied at all the best (in view of what the speaker of c desires at w) admissible ϕ-worlds.

Two quick notes about the new analysis. First, the accessibility-relation semantics sketched above would have difficulty with the restrictor analysis of conditional antecedents. Suppose the source of the ordering-source (the speaker) desires to vote for Obama in the Democratic primary: she strictly prefers him to any of the other candidates. But she also strictly prefers Clinton to any of the other candidates, save Obama. The following conditional seems false of her.

(50) If I don't vote Obama, I ought to vote Gravel. [$\approx O(\neg$*vote-O*$)($*vote-K*$)$]

There are no worlds compatible with the speaker's desires where she does not vote Obama, so the accessibility semantics predicts $O(\neg$*vote-O*$)($*vote-K*$)$ vacuously true.[29] Because an ordering-source semantics selects best worlds (worlds satisfying enough desires) rather than ideal worlds (worlds satisfying every desire) for the domain of the modal, it avoids this prediction.

Second, axiomatizing (constraining the class of models for) the logic is easy enough: we replace **K** and **D** with the following axioms, respectively, and stipulate that they are valid in the class of all models for $\mathcal{L}_{DLK}$.

(51) $O(\pi)(\phi \to \psi) \to (O(\pi)(\phi) \to O(\pi)(\psi))$

(52) $O(\phi)(\psi) \to \neg O(\phi)(\neg\psi)$

Making use of this logic requires revising the Basic Propositional Imperative Language. We replace the imperative clause in the definition of $\mathcal{L}_{PI}$ with (53) and term the new language $\mathcal{L}_{PIK}$.

(53) If $\phi, \psi \in \mathcal{L}_P$, then $!(\phi)(\psi) \in \mathcal{L}_{PIK}$

28. Obviously, then, a one-place deontic necessity operator is $\mathcal{L}_{DLK}$-definable, in terms $O(\cdot)(\cdot)$.

29. In point of fact, **D** forbids vacuous truth of deontic necessities. But it is not clear how to avoid it in the accessibility relation semantics we have stated.

Finally, following roughly the direction sketched in the analysis of conditional imperatives of Schwager (2006), we replace the non-Boolean clause of the semantics in (36).[30]

(54) If $\phi = !(\psi)(\pi)$ (for ψ, $\pi \in \mathcal{L}_P$):
$\mathcal{M}, c, w \Vdash_{\mathcal{L}_{PIK}} \phi$ iff $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLK}} O(\psi)(\pi)$

A restricted imperative $!(\phi)(\psi)$ is required/in force in $\mathcal{M}$ at $\langle c, w\rangle$ just in case ψ is satisfied at all the best admissible ϕ-worlds. A conditional imperative (*if* ϕ)(*stit* ψ) will be generally analyzed as a restricted $\mathcal{L}_{PIK}$ imperative—$!(\phi)(\psi)$—while a bare imperative *stit* ϕ will be analyzed as a vacuously restricted imperative—$!(\top)(\phi)$.

Comparison with narrow-scoping. An immediate advantage of this approach over narrow-scoping is obvious: the posited logical forms for conditional imperatives are directly interpretable (i.e., by us) in terms of their requirement conditions. We may read $!(\phi)(\psi)$ as expressing the restricted command *In all the best* ϕ *possibilities*, *stit* ψ.

There are also, however, illusory advantages, e.g., those claimed by Schwager (2006). Schwager claims that this sort of modal semantics is uniquely well-situated to handling conditional imperatives with overt quantificational material in consequent position. Consider, for example, (55a), the salient reading of which is indicated in (55b):

(55) a. If I need aid, always give it to me!
b. *stit*($\forall w \in \{w \mid$ I need aid at $w\}$: you give me aid at w)

Modeling such constructions is not difficult, given a suitable extension of $\mathcal{L}_{PIK}$. We would introduce into the language a two-place (i.e., domain-restrictable) necessity modal $\alpha(\cdot)(\cdot)$ corresponding to the denotation of the quantificational adverbial *always* (cf. Lewis 1975) and represent (55a) with the following logical form.

(56) $!(\top)(\alpha(\phi)(\psi))$

Schwager takes this to be a problem for narrow-scope analyses—and, more generally, any so-called "hypothetical speech-act" analysis which treats a conditional imperative (*if* ϕ)(*stit* ψ) as generating an unconditional imperative *stit* ψ when the antecedent information ϕ is, in some sense, available. Any such analysis will, Schwager argues, get (55a) wrong. To illustrate, narrow-scope analyses are supposedly committed to representing (55a) with the following logical form.

(57) $\phi \rightarrow !\alpha(\top)(\psi)$

This formula will be required in $\mathcal{M}$ at $\langle c, w\rangle$ just in case ϕ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$ or $!\alpha(\top)(\psi)$ is required in $\mathcal{M}$ at $\langle c, w\rangle$. Supposing that ϕ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$, it follows that $!\alpha(\top)(\psi)$ is required there. But this is surely wrong: supposing I need aid, the content of my command in (55a) does not require that you *always give me aid*.

But this is not convincing. Schwager is implicitly supposing that the narrow-scoper must treat the imperative in (55a) as having the surface form in (58a). In fact, she could

30. While fairly close to the analysis in Schwager (2006), our approach differs materially in two respects. One, Schwager thinks of natural language imperative operators as *literally* deontic necessity modals (with exclusively performative interpretations); we do not. Two, Schwager identifies the modal base with a Stalnakerian Common Ground (Stalnaker 1978, 2002). Although we regard the modal base as informational, we are agnostic about how best to characterize it.

(indeed should) treat it as having the surface form in (58b).

(58) a. $(\textit{if }\phi)(\textit{stit always }\psi)$
b. $\textit{stit}((\textit{if }\phi)(\textit{always }\psi))\ [\approx\ !\alpha(\phi)(\psi)]$

Nothing blocks the narrow-scoper from (i) distinguishing "genuine" conditional imperatives (those treated with surface form (58a)) from "pseudo" conditional imperatives (those treated with surface form (58b)); and (ii) holding that only genuine conditional imperatives are covered by her analysis. We will have to content ourselves with a merely conceptual advantage over narrow-scoping.

Comparison with wide-scoping. Whether or not we endorsed the imperative analogue of modus ponens, we saw that wide-scope analyses failed. But a wide-scoper might reasonably think that a restrictable imperative operator does no better on this score. Choose any $\mathcal{M} = \langle D, W, \mathcal{C}, V\rangle$, and suppose, as before, that ϕ_1 and ϕ_2 partition the set of worlds compatible with the modal base f_c at arbitrary $w \in W$. Suppose additionally that $O(\phi_1)(\psi_1)$ $[\approx (\textit{if }\phi_1)(\textit{stit }\psi_1)]$ and $O(\phi_2)(\psi_2)$ $[\approx (\textit{if }\phi_2)(\textit{stit }\psi_2)]$ are satisfied in $\mathcal{M}$ at $\langle c, w\rangle$. It follows that $\mathcal{M}, c, w \vDash O(\psi_1 \vee \psi_2)$.

Proof. Choose $v \in \bigcap f_c(w)$. In view of Monotonicity (32b):

- If $v \in \bigcap(f_c(w) \cup [\![\phi_1]\!]_{\mathcal{M},c})$ and $v \in \textit{min}(f_c(w), \leq_{g_c(w)})$, then $v \in \textit{min}(f_c(w) \cup [\![\phi_1]\!]_{\mathcal{M},c}, \leq_{g_c(w)})$
- If $v \in \bigcap(f_c(w) \cup [\![\phi_2]\!]_{\mathcal{M},c})$ and $v \in \textit{min}(f_c(w), \leq_{g_c(w)})$, then $v \in \textit{min}(f_c(w) \cup [\![\phi_2]\!]_{\mathcal{M},c}, \leq_{g_c(w)})$

Since ϕ_1 and ϕ_2 partition $\bigcap f_c(w)$, either $v \in \bigcap(f_c(w) \cup [\![\phi_1]\!]_{\mathcal{M},c})$ or $v \in \bigcap(f_c(w) \cup [\![\phi_2]\!]_{\mathcal{M},c})$. It follows that:

- $\textit{min}(f_c(w), \leq_{g_c(w)}) \subseteq$ $\textit{min}(f_c(w) \cup [\![\phi_1]\!]_{\mathcal{M},c}, \leq_{g_c(w)}) \cup \textit{min}(f_c(w) \cup [\![\phi_2]\!]_{\mathcal{M},c}, \leq_{g_c(w)})$

Since $O(\phi_1)(\psi_1)$ and $O(\phi_2)(\psi_2)$ are satisfied in $\mathcal{M}$ at $\langle c, w\rangle$, we have:

- $\textit{min}(f_c(w) \cup [\![\phi_1]\!]_{\mathcal{M},c}, \leq_{g_c(w)}) \subseteq [\![\psi_1]\!]_{\mathcal{M},c}$
- $\textit{min}(f_c(w) \cup [\![\phi_2]\!]_{\mathcal{M},c}, \leq_{g_c(w)}) \subseteq [\![\psi_2]\!]_{\mathcal{M},c}$

These facts together yield: $\textit{min}(f_c(w), \leq_{g_c(w)}) \subseteq [\![\psi_1]\!]_{\mathcal{M},c} \cup [\![\psi_2]\!]_{\mathcal{M},c}$. Then, by the semantics for $O(\cdot)(\cdot)$ of (49), $O(\psi_1 \vee \psi_2)$ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$.

Then $!(\psi_1 \vee \psi_2)$ is required $\mathcal{M}$ at $\langle c, w\rangle$. This is precisely what we took to doom the wide-scope analysis. No account of these sorts of conditionals on the market—e.g., the shifty conditionals of Gillies (2008); Kolodny & MacFarlane (2008)—happens to fare any better, so long as the semantics assumes Monotonicity.

3.4.5 Against Monotonicity

We are not wedded to Monotonocity. Indeed, insofar as there is a rational obligation on agents to be expected utility maximizers, it seems plausible that we *should not* be. The motivating intuition here is that, in ranking possibilities, an agent should, other things equal, privilege those desires which she can reliably expect to be able to fulfill over those desires which she cannot. The more information available to an agent, the *more desires she can reliably expect to fulfill*, and the *stricter* the criteria for remaining at the top of the ranking. While it is true that a world's competition in the ranking shrinks with the addition of information, information acquisition also makes it harder for a world to meet the demands of the privileged desires. The tension with Monotonicity is obvious.

As noted above, Monotonicity (in our sense) follows directly from the monotonicity (order-preservingness) with respect to $\leq_{g_c(w)}$ of a partial identity function mapping from a set of worlds into one of its subsets. While we will want to preserve this lowercase monotonicity—monotonicity of a partial identity function with respect to the preorder we use to rank worlds—we can still avoid uppercase Monotonicity by letting the ordering-source vary according to the information contained in a body of relevant information. We will attempt a sketch of a semantics that implements this idea here.

Suppose we have a modal metalanguage built on top of the Boolean propositional language $\mathcal{L}_P$ in which $\delta\phi$ is a term of the language (when $\phi \in \mathcal{L}_P$), roughly to be read as *the (action of) seeing to it that* ϕ, and $\Omega\delta\phi$ is a formula of the language, roughly to be read as *the (action of) seeing to it that* ϕ *occurs*. We will save rigorous development of a modal language of action and its semantics for our discussion of Propositional Dynamic Logic. We only use it now to gesture at a way of jettisoning Monotonicity from an ordering-source semantics.

The intuitive idea is to select the best worlds relative to both an ordering source *and relevant information*. Let $\mathcal{P} \subseteq 2^W$ be a set of propositions.

(59) *ch*$(g_c(w), \mathcal{P})$ is the set of propositions p such that, for some ϕ and $\delta\psi$:
$[\![\phi]\!] = p$
$\bigcap \mathcal{P} \subseteq [\![\Omega\delta\psi \rightarrow \phi]\!]$
$\bigcap \mathcal{P} \cap [\![\Omega\delta\psi]\!] \neq \emptyset$

To unpack: *ch*$(g_c(w), \mathcal{P})$ gives the set of propositions in $g_c(w)$ that are known (by the lights of the information in $\mathcal{P}$) to be fulfillable by a possible (by the lights of $\mathcal{P}$) action.[31] In the terminology of Kolodny & MacFarlane (2008), *ch*$(g_c(w), \mathcal{P})$ gives the set of *choices* relative to an ordering source and body of information.[32] We next define a preorder according to the choices relative to an ordering source and body of information.

(60) $v \leq_{ch(g_c(w),\mathcal{P})} u$ iff
$\{p \in ch(g_c(w), \mathcal{P}) \mid u \in p\} \subseteq \{p \in ch(g_c(w), \mathcal{P}) \mid v \in p\}$

Correspondingly, we redefine the set of best worlds relative to a body of information $\mathcal{P}$ and ordering-source $g_c(w)$—*sel*$(\mathcal{P}, g_c(w))$—as follows. Note that worlds are ordered according to the choices relative to $g_c(w)$ and $\mathcal{P}$—i.e., the desires (members of $g_c(w)$) that are known (by the lights of $\mathcal{P}$) to be satisfiable by a definite course of action.

(61) *sel*$(\mathcal{P}, g_c(w)) =$
$\{v \in \bigcap \mathcal{P} \mid \forall u \in \bigcap \mathcal{P} : u \leq_{ch(g_c(w),\mathcal{P})} v \Rightarrow v \leq_{ch(g_c(w),\mathcal{P})} u\}$

We have at this point shaken the Monotonicity property. It is clearly not the case that a world that counts as best relative to a set of desires $g_c(w)$ and body of information $\mathcal{P}$ will count as best relative to the same set of desires and a more specific body of information

31. It would be interesting to implement this sort of idea probabilistically, so that desires could be ranked (rather than simply admitted or eliminated) by an agent's assessment of the likelihood that they will be fulfilled. Although Yalcin (2007) provides some relevant background, the use of probabilistic methods in formal semantics is fairly uncharted territory. I leave this to a future project.

32. It is their terminology, but they do not implement the idea in terms of a formal ordering-source semantics for deontic modals, as we do here.

$\mathcal{P}'$. A desire that does not count as a choice relative to a comparatively unspecific body of information will often count as a choice relative to a comparatively specific body of information, and it is choices, rather than mere desires, that are used to rank worlds.

To finish the implementation, we rewrite the satisfaction clause for the two-place deontic operator $O(\cdot)(\cdot)$ as follows.

(62) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLK}} O(\phi)(\psi)$ iff
$\forall v \in \mathit{sel}(f_c(w) \cup [\![\phi]\!]_{\mathcal{M},c}, g_c(w)) : \mathcal{M}, c, v \vDash_{\mathcal{L}_{DLK}} \phi$

On this semantics, $O(\phi)(\psi)$ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$ iff ψ is satisfied at the best—according to the choices relative to $f_c(w) \cup [\![\phi]\!]_{\mathcal{M},c}$ and $g_c(w)$—ϕ-worlds compatible with $f_c(w)$. The requirement conditions for two-place imperatives are left unchanged.

This is enough to block the proof given at the beginning of this section. Since the apparatus here is very abstract, we will try to make it intuitive by applying it informally to the problem case of Kolodny & MacFarlane (2008). Suppose the speaker desires that all ten miners are saved. This desire does not count *as a choice* relative to the information available in the basic context, since that information, by supposition, does not settle which shaft the miners in, and, so, does not settle any definite way of achieving this desire. And so this desire cannot be used to privilege worlds in which all ten miners are saved. So we will not generally expect either $O(\top)(\mathit{bl\_A}) \vee O(\top)(\mathit{bl\_B})$ [$\approx$ *ought bl_A* $\vee$ *ought bl_B*)] or $O(\top)(\mathit{bl\_A} \vee \mathit{bl\_B})$ [$\approx$ *ought*($\mathit{bl\_A} \vee \mathit{bl\_B}$)] to come out true with respect to the basic context. The desire does, however, count as a choice relative to the modal base augmented with the proposition that the miners are in shaft A, and so we will generally expect $O(\mathit{in\_A})(\mathit{bl\_A})$ [$\approx$ (*if in_A*)(*ought bl_A*)] to come out true.

This semantics saves the two-place analysis of the imperative operator from immediate empirical difficulty by avoiding Monotonicity. This means, inter alia, that it cannot help the wide-scoper. As the reader may easily verify, all of the wide-scoper's problems stem from foundational facts about the semantics of deontic logic, and none bear any connection whatever to Monotonicity.

This will for the most part conclude our explicit discussion of conditional imperatives in this paper, and we will simply take the two-place analysis of imperative and deontic operators for granted throughout the rest of it. Although conditional imperatives will no longer be a major topic of interest, the reader should be aware that by placing our focus on two-place operators, we have essentially placed our focus on conditional imperatives, and are treating unconditional imperatives as a special case. We turn our attention now to possible treatments of the Ross Paradox in terms of deontic modal logic.

3.5 The Ross Paradox

The Ross Paradox may be thought to imperil any analysis of imperative logic in terms of deontic modal logic, for two reasons. First, the following conditional is valid in the class of all models for deontic modal logic.

(63) $O\phi \rightarrow O(\phi \vee \psi)$

Second, it might be thought that this is the *right* result: if it ought to be that you post the letter, then it ought to be that you post or burn it, i.e., *It ought to be that you post the letter* does indeed entail *It ought to be that you post or burn the letter*. While it is certainly *possible* to hear *It ought to be that you post or burn the letter* as expressing a free choice permission to burn the letter, the free choice reading is demonstrably less salient than in the case of the disjunctive imperative *Post or burn the letter*.[33] If that is right, then the semantics for imperatives we have been developing will need to be modified; it will predict that $!\phi \Vdash !(\phi \vee \psi)$.

There is no really systematic way to evaluate intuitions about such cases. But whether or not the intuitions are correct, we will show that there is no argument here, as such, against treating imperative logic in terms of deontic logic. In this section, I sketch two possible lines of response. The first—inspired by the so-called Montague-Scott or Neighborhood Semantics treatment of the problem of logical omniscience—involves rejecting the intuitions, maintaining the analysis of ! in terms of O, and revising the semantics for deontic modal logic from the ground up, so that $O\phi \rightarrow O(\phi \vee \psi)$—or, more accurately, $O(\pi)(\phi) \rightarrow O(\pi)(\phi \vee \psi)$—is *not* valid in the class of all models for the deontic modal logic. The second involves accepting the intuitions, but adapting to them by formulating a new analysis on which the imperative operator ! is not treated strictly in terms of the deontic necessity operator O, but also in terms of the deontic permission operator P.

3.5.1 Neighborhood Semantics

Overview. Neighborhood semantics (a.k.a. Montague-Scott semantics[34]) is a non-relational generalization of the relational (Kripke) semantics for modal languages. Neighborhood models are standardly defined as follows.

(64) $\mathcal{M} = \langle W, N, V \rangle$
$N : W \mapsto 2^{2^W}$

N is a function from worlds to sets of propositions, where $N(w)$ roughly yields the set of propositions that are necessary (in whatever sense of necessity we are interested in using the logic to model) at w. Neighborhood semantics alters the relational semantics for a generic modal formula $\Box\phi$ as follows. Let $\mathcal{M} = \langle W, N, V \rangle$ be a neighborhood model.

(65) $\mathcal{M}, w \vDash \Box\phi$ iff $[\![\phi]\!]_{\mathcal{M}} \in N(w)$

As a function from worlds to sets of propositions, N is a good deal like the conversational backgrounds (modal bases and ordering-sources) of which we have been making use throughout the paper. But its role in the semantics is quite different: rather than being treated as a universal over accessible (best) worlds, the necessity modal simply checks to see if the proposition expressed by its prejacent is a member of the set of propositions necessary at a world. Because closure conditions on N are entirely flexible, neighborhood semantics has an easier time modeling phenomena like the non-closure of knowledge or belief under logical consequence than the relational semantics.

Application. We have at the ready a natural candidate to play the role of the neighborhood function in our semantics—the erstwhile ordering-source. Let $\mathcal{M} = \langle D, W, \mathcal{C}, V \rangle$. The obvious way of extending neighborhood semantics to the basic deontic language is as

33. Interestingly, changing the modal to *may* seems to reverse some people's intuitions (and theoretical opinions). Aloni (2007), for example, gives a semantic account of free choice permissions in cases where *may* scopes over a disjunction, but a pragmatic account in cases where *must* scopes over a disjunction. The asymmetry is unappealing, but I won't try to directly resist it.

34. See Chellas (1980) for relevant references and background. The idea for using neighborhood semantics to handle the Ross Paradox originates with Segerberg (1990), and the analysis presented here is heavily indebted to that paper.

follows:

(66) $\mathcal{M}, c, w \vDash O\phi$ iff $[\![\phi]\!]_{\mathcal{M},c} \in g_c(w)$

How might we extend this basic idea to the more complicated deontic language $\mathcal{L}_{DLK}$? The following semantics naturally suggests itself.

(67) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLK}} O(\phi)(\psi)$ iff
$\forall v \in \mathit{sel}(f_c(w) \cup [\![\phi]\!]_{\mathcal{M},c}, g_c(w)) : [\![\psi]\!]_{\mathcal{M},c} \in g_c(v)$

Informally, $O(\phi)(\psi)$ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$ iff ψ is desired at the best—according to the choices relative to $f_c(w) \cup [\![\phi]\!]_{\mathcal{M},c}$ and $g_c(w)$—ϕ-worlds compatible with $f_c(w)$. On this semantics, the ordering-source plays a dual role: while it still functions as an ordering-source role on worlds, in addition, it *directly* tells us what sorts of things are required with respect to every world. The ordering-source plays, in the terminology of Segerberg (1990), the role of a "command system"—a semantic device designed "to keep track of the commands issued [or, in our setup, issuable] by the authority" Segerberg (1990: 204).

Some notation will make our life easier, while also bringing out the parallels with neighborhood semantics as traditionally conceived. We define a "genuine" neighborhood function $G_{\mathcal{M},c} : W \mapsto 2^{2^W \times 2^W}$ as follows:

(68) $G_{\mathcal{M},c} = \lambda w.\{\langle [\![\phi]\!]_{\mathcal{M},c}, [\![\psi]\!]_{\mathcal{M},c}\rangle \mid \mathcal{M}, c, w \vDash_{\mathcal{L}_{DLK}} O(\phi)(\psi)\}$

It follows from this definition that:

(69) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLK}} O(\phi)(\psi)$ iff $\langle [\![\phi]\!]_{\mathcal{M},c}, [\![\psi]\!]_{\mathcal{M},c}\rangle \in G_{\mathcal{M},c}(w)$

The semantic clause for imperative formulae will remain as before. Most of the work will go into axiomatizing the deontic side of the logic in the appropriate way. To that end, we put forward the following minimal list of deontic axioms for $\mathcal{L}_{DLK}$:

(70) $O(\phi)(\psi) \to \neg O(\phi)(\neg\psi)$ [**D**]

(71) $(O(\pi)(\phi) \wedge O(\pi)(\psi)) \to O(\pi)(\phi \wedge \psi)$ [**A**]

D is a minimal consistency requirement on commands. The logician of content will regard the Aggregation axiom **A** as a reasonable "stand-in" for $O(\pi)(\phi \to \psi) \to (O(\pi)(\phi) \to O(\pi)(\psi))$ [**K**].[35] As Segerberg writes about a similar axiom, "This condition reflects the fact that when an authority issues commands, then he or she or it means for them all to be obeyed" (Segerberg 1990: 220). She will, on the other hand, regard neither **K** nor **NEC**—i.e.: $\vDash \phi \Rightarrow \vDash O(\pi)(\phi)$—as desirable additions to the axiomatization.

The case against **K**: suppose it is required that: you either don't post the letter ($\neg\phi$), or you either FedEx or burn it ($\psi \vee \chi$)—i.e., it is required that: $\neg\phi \vee (\psi \vee \chi)$. Suppose it is also required that you do post the letter (and that burning entails failure to post).[36] Should

35. Strengthening **A** to a biconditional **A+** is out of the question, for two reasons:

- Since $\phi \wedge (\neg\phi \vee \psi) \equiv \phi \wedge \psi$, $O(\pi)(\phi \wedge \psi) \equiv O(\pi)(\phi \wedge (\neg\phi \vee \psi))$. Then by **A+**, we have $O(\pi)(\phi \wedge \psi) \vDash O(\pi)(\neg\phi \vee \psi)$, and therefore $!(\pi)(\phi \wedge \psi) \Vdash !(\pi)(\neg\phi \vee \psi)$. This is as problematic as the Ross inference.
- **A+** entails **K**. Proof: suppose **A+** and that both $O(\pi)(\phi \to \psi)$ and $O(\pi)(\phi)$ are satisfied in $\mathcal{M}$ at $\langle c, w\rangle$. By **A+**, $O(\pi)((\phi \to \psi) \wedge \phi)$ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$, in which case—since $(\phi \to \psi) \wedge \phi \equiv \phi \wedge \psi$—$O(\pi)(\phi \wedge \psi)$ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$. By **A+**, $O(\pi)(\psi)$ is satisfied in $\mathcal{M}$ at $\langle c, w\rangle$.

36. There is, I admit, something odd about this constellation of requirements. The first requirement ($\neg\phi \vee (\psi \vee \chi)$) seems to grant permissions (to fail to post, and to burn) that the second takes away. But it is not an oddness that the view on offer is able to leverage in a defense of **K**. The view on offer can only explain failure of

it follow that that you should either FedEx or burn it? Not, I should think, for the logician of content pursuing the neighborhood semantics resolution of the Ross Paradox: endorsing *stit*$(\neg\phi \vee (\psi \vee \chi))$ (*Either don't post it, or FedEx or burn it!*) and *stit* ϕ (*Post it!*) intuitively does not commit an agent to endorsing *stit*$(\psi \vee \chi)$ (*FedEx or burn it!*). Endorsing **K** is evidently not consonant with this intuition. The case against **NEC** is straightforward: we will not want to predict commands like *Torture your little brother, or don't* as everywhere required.

Particular axiomatizations of the logic will enforce certain conditions on $G_{\mathcal{M},c}$ (and thereby restrict the class of neighborhood models for the logic). The axioms we have chosen (**D** and **A**) require the following conditions on $G_{\mathcal{M},c}(w)$, respectively. Let $\mathcal{M} = \langle D, W, \mathcal{C}, V\rangle$ and choose $w \in W$, $c \in \mathcal{C}$.

(72) a. If $\langle [\![\phi]\!]_{\mathcal{M},c}, [\![\psi]\!]_{\mathcal{M},c}\rangle \in G_{\mathcal{M},c}(w)$, then $\langle [\![\phi]\!]_{\mathcal{M},c}, [\![\neg\psi]\!]_{\mathcal{M},c}\rangle \notin G_{\mathcal{M},c}(w)$.
b. If $\langle [\![\pi]\!]_{\mathcal{M},c}, [\![\phi]\!]_{\mathcal{M},c}\rangle \in G_{\mathcal{M},c}(w)$ and $\langle [\![\pi]\!]_{\mathcal{M},c}, [\![\psi]\!]_{\mathcal{M},c}\rangle \in G_{\mathcal{M},c}(w)$, then $\langle [\![\pi]\!]_{\mathcal{M},c}, [\![\phi \wedge \psi]\!]_{\mathcal{M},c}\rangle \in G_{\mathcal{M},c}(w)$.

It is evident that we do not, in the axiomatization we have offered, have anything like general closure under logical consequence:

(73) If $\langle p, q\rangle \in G_{\mathcal{M},c}(w)$ and $\langle p, r\rangle \in G_{\mathcal{M},c}(w)$, then $\langle p, s\rangle \in G_{\mathcal{M},c}(w)$, for $s \supseteq q \cap r$.

We have only very circumscribed closure under logical consequence, of the sort required by the Aggregation axiom. This is, of course, sufficient to block the Ross Paradox.

There are, however, drawbacks, the most significant of which is a failure to say anything about the permissive content of imperatives. The *source* of the Ross Paradox is the felt free choice permissions of a disjunctive imperative. Neighborhood semantics for imperatives is motivated by this intuition, but tries to do it justice by circumscribing closure conditions on the neighborhood function. It is not clear what explanans has to do with explanandum. Or, to put it differently, if circumscription of closure conditions is warranted, then it is warranted in virtue of the permissive content of disjunctive imperatives. A logic in which that dimension of imperative content is explicitly represented would provide a more satisfying account of the intuitions behind the Ross Paradox than one in which it is not.

The failure to explicitly represent permissive content manifests in predictive gaps. Recall the pair (21a) and (21b), repeated here, which we took tenatively to motivate a semantic resolution of the Ross Paradox:

(74) a. ?Post or burn the letter. But you may not burn it!
b. ?Post or burn the letter. But do not burn it!

If this data motivates a semantic resolution of the Ross Paradox, then such a resolution ought to have something to say by way of explaining the data. The neighborhood semantics resolution of the Ross Paradox explains only the failure of $!(\pi)(\phi)$ to imply $!(\pi)(\phi \vee \psi)$, and does so by appeal to lack of closure of $\mathcal{G}_{\mathcal{M},c}$ under logical consequence (cf. 73). The absence of a closure property cannot be leveraged to explain the oddness of "synchronic" cancellation constructions. Nor can it be leveraged to explain the oddness of "diachronic" cancellation constructions, of the sort we used to argue against axiomatizing the neighborhood semantics with **K**.

Worse, losing **K** means automatically committing ourselves to a logic of the sort we

---

entailments, by appeal to lack of certain conditions on the neighborhood function. Because permissive content is not built into the semantics, the view has nothing to say about inconsistencies in permissive and command content.

saw Castaneda arguing for in §2.3 (recall argument pattern 17). This means not generally being able to draw an imperative conclusion which is weaker (in terms of what it requires) from imperative premises which are together stronger.[37] While this is certainly a reasonable logic, there is also, we argued, a reasonable logic of content which counts the argument in (17) *valid*. We are thus faced with a dilemma: either endorse **K** (in which case, in the context of a neighborhood semantics, we predict a close relative of the Ross inference valid), or reject **K** (in which case, we are unable to characterize a logic in which the argument in (17) is valid). Both horns are rather unappealing.

We now move on to consider a treatment of the Ross Paradox that is conservative with respect to the semantics for deontic logic we have been developing, on which the relevant work is done by positing explicit permissive content, rather than by fiddling with the semantics for the O operator.

3.5.2 Permission Analyses

Understanding imperatives as having content along two dimensions—a requiring dimension (its command content) and a permitting dimension (its permission content)—is a useful way of thinking about imperative content, particularly in connection with the Ross Paradox. So-called permission analyses of the disjunctive imperative $!(\phi \vee \psi)$ are distinguished by (i) endorsing such a two-dimensional analysis of imperative content and (ii) holding that the permissive content of $!(\phi \vee \psi)$, on one reading, expresses permissions both to see to it that ϕ and to see to it that ψ.

The most rigorous development of a permission analysis is found in an excellent recent paper by Maria Aloni (Aloni 2007). Aloni's idea is to understand the imperative operator in terms of a single, complex modal operator that expresses requirement of $\phi \vee \psi$ and free-choice permission to fulfill this requirement by securing the truth of either disjunct, which we represent schematically as *may*$_F(\phi \vee \psi)$. The idea that imperatives express free-choice permissions, by itself, only gets us a little way to a solution, however, since the paradox of free-choice permission—very roughly, the problem of how to explain why *may*$_F(\phi \vee \psi)$, on a salient interpretation, implies both *may* ϕ and *may* ψ—is, as already noted, a vexing puzzle in its own right.

Alternative Semantics. The key to Aloni's analysis is an idea borrowed from recent work on the semantics and pragmatics of questions (see especially Aloni & van Rooy 2002). We understand disjunctions as being associated with ("inducing") alternatives. Alternatives in turn serve as the objects of higher-order intensional or dynamic operators. In the case of questions, we have a dynamic question operator which adds alternatives to a list of topics under consideration ("at issue") in a dialogue; roughly, if ϕ is on the list of topics, then whether or not ϕ is an issue of interest in the dialogue. A polar (yes/no) question $?(\phi \vee \psi)^{\sharp}$[38] fails to induce genuine alternatives (which we model as a singleton alternative set $\{\phi \vee \psi\}$, the contents of which the question operator adds to the conversational topic list). A choice-presenting question $?(\phi^{\sharp} \vee \psi^{\sharp})$ presents both ϕ and ψ as issues of interest; its alternative set is given as $\{\phi, \psi\}$. An illustration: (75) is ambiguous between (76a)

37. In this respect, the resulting logic will bear some similarity to that of Kenny (1966) (see §2.4). This is not to say the logic will be, like Kenny's, nonmonotonic. The neighborhood logic requires that an imperative conclusion be at least as strong as some non-empty subset of premise imperatives, and this will make the corresponding entailment relation monotonic. Kenny's logic requires that an imperative conclusion be at least as strong as the *entire* premise set; this is what is responsible for its nonmonotonicity.

38. Superscripts indicate placement of focal stress. On the semantic relationship between focus and salient alternatives in such linguistic phenomena as ellipsis resolution and presupposition see, e.g., Charlow (2008c); Rooth (1992).

(polar) and (77a) (choice-presenting).

(75) Do you like Rori or Uni?

(76) a. (...Rori or Uni)$^\sharp$? [$\approx ?(\phi \vee \psi)^\sharp$]
b. Yes/no.

(77) a. ...(Rori)$^\sharp$ or (Uni)$^\sharp$? [$\approx ?(\phi^\sharp \vee \psi^\sharp)$]
b. *Yes/no.

When stress is unfocused (76a), a single issue (whether $\phi \vee \psi$) is under discussion, and a *yes/no* answer is appropriate. When stress is focused (77a), there are two topics of interest (whether ϕ, whether ψ). Without elaboration: *yes/no* fails to answer the question.

Aloni (2007) suggests analyzing the imperative operator as a bi-dimensional requirement/ permission modal over alternatives induced by its complement. Informally, the alternatives induced by the complement of the imperative operator represent permitted *ways of complying* with the command (Aloni 2007: 87). Imperatives whose complements induce genuine alternatives expressly permit multiple modes of compliance. Choice offering readings of disjunctive imperatives are associated with non-singleton alternative sets; no-choice readings are not.

Aloni's actual implementation is complex (and, we will see, may be significantly simplified). She introduces a higher-order extension of a first-order deontic language—one distinguished by its use of propositional quantification—to serve as the imperative object language, in which logical forms are given. The logical form of a genuine alternative-inducing disjunction is given as $\exists p(p \wedge (p = \phi \vee p = \psi))$ (to be read: something holds, either ϕ or ψ) and that of a non-alternative-inducing disjunction as $\exists p(p \wedge p = \phi \vee \psi)$ (to be read: something holds, namely $\phi \vee \psi$). There are various methods of generating a set of alternatives, *alt*(π) (where π is a formula in the object language), from such logical forms,[39] but they will not be our concern here. Free choice readings of disjunctive imperatives are analyzed as in (78a), no-choice readings as in (78b). The imperative operator is defined in terms of modals and alternative sets (79).

(78) Analysis of imperatives.
a. $!\exists p(p \wedge (p = \phi \vee p = \psi))$
b. $!\exists p(p \wedge p = \phi \vee \psi)$

(79) Modal reduction.
a. $!\phi := \nabla \mathit{alt}(\phi)$
b. $\nabla(\phi_1, ..., \phi_n) := [P](\phi_1, ..., \phi_n) \wedge O(\phi_1 \vee ... \vee \phi_n)$
c. $[P](\phi_1, ..., \phi_n) := P\phi_1 \wedge ... \wedge P\phi_n$

(80) Definitional identities.
a. $!\exists p(p \wedge (p = \phi \vee p = \psi)) := (P\phi \wedge P\psi) \wedge O(\phi \vee \psi)$
b. $!\exists p(p \wedge p = \phi \vee \psi) := P(\phi \vee \psi) \wedge O(\phi \vee \psi)$

A free choice permission sentence $\mathit{may}_F(\phi \vee \psi)$ is analyzed with the logical form $[P]\mathit{alt}(\exists p(p \wedge (p = \phi \vee p = \psi)))$, i.e., $P\phi \wedge P\psi$. Free choice disjunctive imperatives are thus analyzed in terms of free choice permissions *and* no-chice requirements: they express that either disjunct is permitted, at least one required.[40] We have already discussed why this blocks

39. For a baroque method, see Aloni (2007: 72-5). For something a bit less baroque (but also less general) that relies only on tweaking of the assignment function for propositional variables, see Charlow (2008a).

40. For epistemic *may*, Aloni's proposal is equivalent to the much-discussed proposal of Zimmermann (2000), which analyses disjunctions as *conjunctions of epistemic possibilities*, rendering $\mathit{may}_F(\phi \vee \psi)$ as $\Diamond_e(\Diamond_e\phi \wedge \Diamond_e\psi)$. Where Aloni's analysis comes into its own is disjunctions scoped under deontic *may*—something for which Zimmermann has no story.

the Ross Paradox.

Integrating the two-place treatment of deontic and imperative operators is a matter of construing ∇ as having a restriction argument, and allowing the restriction to percolate downward.

(81) a. $!(\phi)(\psi) := \nabla(\phi)(alt(\psi))$
b. $\nabla(\phi)(\psi_1,...,\psi_n) := [P](\phi)(\psi_1,...,\psi_n) \wedge O(\phi)(\psi_1 \vee ... \vee \psi_n)$
c. $[P](\phi)(\psi_1,...,\psi_n) := P(\phi)(\psi_1) \wedge ... \wedge P(\phi)(\psi_n)$

Simplification. Given that there is no real explanation of *when* the use of (78a) in a formal representation of a disjunctive imperative is preferable to (78b), it is not quite clear what Aloni's complication of the object language ultimately accomplishes. We might instead try an account that left the object language as before, while "associating" disjunctive imperatives directly with one or other of their proposed deontic logical forms. This would have to be done in a particular way. We do not want, for example, to associate formulae of $\mathcal{L}_{PIK}$ with *multiple logical forms* in $\mathcal{L}_{DLK}$—one for a free choice or no-choice interpretation, as the case may be—for two reasons. First, and more critically, we do not wish to make the interpretation of an imperative formula in $\mathcal{L}_{PIK}$ indeterminate. This is exactly what association with multiple logical forms will accomplish. One tangible benefit of Aloni's analysis is the banishment of ambiguity from the "imperative" object language. Second, we do not want to give imperative formulae non-imperative logical forms at all. Modal analyses of the imperative do not have to be eliminativist.[41] We can give requirement conditions for imperative formulae *in terms of* satisfaction-conditions for deontic formulae, without holding that, for example, the imperative operator is literally any species of deontic operator (cf. §3.3).

We will present an analysis in a similar spirit as Aloni's that avoids recourse to an opulent object language, while avoiding both interpretative indeterminacy and eliminativism about imperative logic. The crucial work is done by assuming that context determines salient alternatives for a disjunctive imperative. We will not complicate the context parameter any further, although a realistic treatment of how salient alternative sets are generated for disjunctive imperatives in natural language would have to pay attention to features of context that we are not attempting to represent in our model theory, focus values especially. Instead, we look to the ordering-source and modal base to determine salient alternatives. Disjunctive imperatives receive free choice interpretations by default; no choice interpretations are invoked only in contexts where the modal base provides decisive information about what is permissible in view of the desires of the speaker.

For technical reasons having to do with the mechanics of the two-place imperative operator, we define salient alternatives for imperative formulae of $\mathcal{L}_{PIK}$, rather than formulae of $\mathcal{L}_P$, doing so inductively, as follows.

(82) a. $alt_{\mathcal{M},c,w}[!(\pi)(p)] = p$
b. $alt_{\mathcal{M},c,w}[!(\pi)(\neg\phi)] = \neg\phi$
c. Given $\phi = !(\pi)(\psi_1 \vee ... \vee \psi_n)$, let $opt_{\mathcal{M},c,w}(\phi) =$
$\{\psi_i \mid \bigcap f_c(w) \not\subseteq [\![P(\pi)(\psi_i)]\!]_{\mathcal{M},c}$ and $\bigcap f_c(w) \not\subseteq [\![\neg P(\pi)(\psi_i)]\!]_{\mathcal{M},c}\}$
If $opt_{\mathcal{M},c,w}(\phi) \neq \emptyset$, $alt_{\mathcal{M},c,w}(\phi) = opt_{\mathcal{M},c,w}(\phi)$
Otherwise, $alt_{\mathcal{M},c,w}(\phi) = \phi \vee \psi$

41. Strictly speaking, there are two ways for an imperative logic to be eliminativist. One, by defining imperative formulae as non-imperative formulae. Two, by identifying satisfaction-conditions for imperative formulae with those of non-imperative formulae. Aloni's account is eliminativist in both ways, as is Schwager's. Avoiding the incredulous stare seems to require both an autonomous imperative object language and denying that imperatives have satisfaction-conditions at all.

Genuine alternatives are induced by default by a disjunctive imperative $!(\pi)(\phi \vee \psi)$, absent information that settles the question of the permissibility of ϕ or ψ (if π) *independently*. This is a crude way of generating salient alternatives, but it will do for our purposes. We finish the simplification by revising the requirement conditions for imperatives to incorporate salient alternatives.

(83) If $\phi = !(\pi)(\psi)$ (for $\pi, \psi \in \mathcal{L}_P$):
$\mathcal{M}, c, w \Vdash_{\mathcal{L}_{PIK}} \phi$ iff $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLK}} \nabla(\pi)(alt_{\mathcal{M},c,w}(\phi))$

Simplicity has its virtues. This account is no less explanatory than Aloni's, which explains the difference between free choice and no-choice interpretations of disjunctive imperatives by appeal to properties of their salient alternatives, without giving any account of when certain salient alternatives are generated over others. This account does at least as well—indeed, better, insofar as we provide a precise (if crude) account of how alternative sets are generated—and meanwhile avoids purposeless complication of the object language. We also avoid the above-mentioned indeterminacy of interpretation with respect to $\mathcal{L}_{PIK}$: no difference in logical form is associated with free choice and no-choice interpretations, and both are handled with a single clause of the semantics. Requirement conditions for disjunctive imperatives will depend on facts about context, but we have been committed to this from the beginning.

It is worth noting in closing that our differences with Aloni are not merely methodological. On account of her modal eliminativism about imperative logic, Aloni does not elect (as we do) to utilize an imperative object language in which requirement conditions for disjunctive imperatives can vary with facts about context. Because of this, she (along with any modal eliminativist who endorses a permission analysis of the Ross Paradox) cannot do full justice to the Ross Paradox. Consider the paradox in its original form.

(84) $!\phi \nVdash !(\phi \vee \psi)$

The account we have given predicts something very close in form to this:

(85) $!(\phi)(\psi) \nVdash !(\phi)(\psi \vee \pi)$

We have defined imperative entailment in terms of preservation of modal satisfaction in all models and at all context-world pairs. Because there are contexts in which the modal base does not provide decisive information about what is permissible, there are obviously models $\mathcal{M}$ and context-world pairs $\langle c, w\rangle$ such that $\mathcal{M}, c, w \vDash \nabla(\phi)(\psi)$, but $\mathcal{M}, c, w \nvDash \nabla(\phi)(\psi \vee \pi)$.

The best a modal eliminativist can do is say that the $\mathcal{L}_{DLK}$ logical form associated with a free choice interpretation of $!(\phi)(\psi \vee \pi)$ is not entailed by the $\mathcal{L}_{DLK}$ logical form of $!(\phi)(\psi)$. (But, of course, we *knew* that already, de re anyway.) The Ross Paradox is explained away as a side-effect of coarse-grainedness in the formal imperative language—its inability to effectively represent the differences between the relevant logical forms. Insofar, then, as we are inclined not to endorse an error theory about the Ross Paradox, we will be inclined to reject the modal eliminativist's analysis of it.

How do we do with the problems for the neighborhood semantics account of the Ross Paradox? The oddness of synchronic cancellation constructions is explained by appeal to the default permissive content of the disjunctive imperative. So long as *Post or burn the letter!* (on the default interpretation) is required in $\mathcal{M}$ at $\langle c, w\rangle$, it cannot be the case that *Do not burn the letter!* is also required in $\mathcal{M}$ at $\langle c, w\rangle$. The permission expressed by the former is simply inconsistent with the command content of the latter. Similarly for diachronic cancellation constructions: *Either don't post the letter, or FedEx or burn it!*

expresses, on the default interpretation, a permission that it is inconsistent with *Post the letter!*. These imperatives cannot both be required in $\mathcal{M}$ at $\langle c, w\rangle$. *FedEx or burn it!* follows from them only in the degenerate sense that everything follows from a contradiction. Insofar as logicians of content do not wish to commit agents who endorse imperatives expressing inconsistent permissions and commands to endorsing everything, they can require that, for an argument form to be valid, there be at least one model and context-world pair in which all the premise imperatives are required.

Note that, on this tack, we still have **K** for deontic O. This means we can, if we please, endorse the validity of (17).[42] What we *cannot* endorse is an imperative analogue of **K**:

(86) If $\mathcal{M}, c, w \Vdash !(\pi)(\phi \to \psi)$ and $\mathcal{M}, c, w \Vdash !(\pi)(\phi)$, then $\mathcal{M}, c, w \Vdash !(\pi)(\psi)$.

But this is a natural consequence of understanding imperative operators as bi-dimensional—$!(\pi)(\psi)$ may express a permission with which $!(\pi)(\phi)$ is inconsistent. In cases where it does not (such as 17), the argument from $!(\pi)(\phi \to \psi)$ and $!(\pi)(\phi)$ to $!(\pi)(\psi)$ is generally predicted good.

3.6 Conclusion

The object of this section was to devise a logic of content for a formal imperative language in terms of a relatively standard model theory for deontic languages. We did this, and showed that, suitably amended, it could be extended to a palatable analysis of conditional imperatives and the Ross Paradox. In the next section, we shift our focus to dynamic logics of planning. The desiderata for such logics remain fairly constant: we will pursue a semantics in terms of the model theory for a deontic modal language, and argue that it too can handle conditional imperatives as well as the Ross Paradox. Shifting to more powerful dynamic and deontic object languages will introduce extra complexity to our project (while also bringing some rewards), but our overarching target—developing a model theory for a formal imperative language that hews closely to the model theory for a deontic language—will remain basically the same.

4 DYNAMIC DEONTIC LOGIC OF ACTION

Dynamic accounts of the imperative come in two distinct flavors. Both are, at first pass, most consonant with the motivations that underlie logics of planning. First, there are accounts of how imperatives govern the planning behavior of agents, where the notion of *government* is understood in a static sense. Imperatives are seen to govern planning behavior via embodying constraints on the way an agent may interact with her surroundings at a given point in time (rather than via effecting changes on the agent's planning behavior). Despite employing a static (i.e., synchronic) notion of government, such accounts are nevertheless dynamic, in the important sense that they view imperatives as constraining the ways agents may effect *transitions* between "states" of the world and are most naturally implemented with a Propositional Dynamic Logic of Action (PDLA), in which (i) we have an object language that is capable of representing actions and of expressing claims about how actions affect the world; (ii) actions are typed as relations between states of the world.

42. A thorny issue: supposing $\phi \to \psi$ abbreviates $\neg\phi \vee \psi$, we have it that the premise imperatives of (17) are default inconsistent. *stit: if you read the book, you come see me* expresses by default a permission (not to read the book) that is inconsistent with *stit: you read the book*. This does not seem right, which pushes us toward a language in which (i) $\phi \to \psi$ does not abbreviate $\neg\phi \vee \psi$, (ii) $\phi \to \psi$ and $\neg\phi \vee \psi$ are truth-functionally equivalent in non-imperative contexts, but (iii) not generally intersubstitutable in imperative contexts, on account of the syntactic sensitivity of the method of salient alternative generation mechanism, which has disjunctions as alternative-presenting by default, implications not.

The PDLA approach will be the focus of this section of the paper.

Second, there are accounts of the imperative that are dynamic in the sense of Groenendijk & Stokhof (1991); Veltman (1996). Such accounts are in the business of modeling the *obligation-imposing* function of imperatives. They also regard imperatives as governing planning behavior, but give the notion of government a dynamic (diachronic) gloss: imperatives govern planning behavior by effecting diachronic changes on an agent's plans, rather than by embodying synchronic constraints on an agent's planning behavior. Such accounts are not necessarily semantic in character—indeed, the most prominent example of such in the linguistics literature style themselves as pragmatic accounts of the context change potentials of imperative utterances. As we shall see, however, there is a natural way of leveraging imperative context change potentials for the sake of formulating a genuine Dynamic Semantics for the Imperative (DSI). We will tackle these issues in the subsequent section of the paper.

As we will also see, and *pace* Portner (2008), there is no reason to see these distinct conceptions of imperative dynamics as standing in any sort of opposition to one another. Accounts in the vein of PDLA model something worth modeling, as do accounts in the vein of DSI. What's more, it's clear that the phenomena they are in the business of modeling—synchronic and diachronic constraints on planning behavior, respectively—are closely interrelated. I take this to motivate the development of formal apparatus that can fill the theoretical needs of both kinds of dynamic account. Trying to do justice to this motivation will be an important theme of the remainder of the paper.

4.1 Language

We begin by complicating our imperative and deontic object languages, so that both are capable of talking about actions and their effects on the world.[43] There are good reasons for doing this, in the context of a logic of planning. The structure of planning is prima facie distinct from the structure of desiring. Desires provide goals. Planning is reasoning about how an agent may *act* to achieve those goals. In representing planning, it is useful, then, to have an object language that is capable of talking about actions. Planning also involves reasoning about the consequences of actions on an agent's future planning: agents often engage in reasoning at a given time t how certain courses of action will introduce constraints on their behavior at a later time t', in view of the various hypothetical constraints on planning behavior that are in force at t. In representing this facet of planning, it is useful to have an object language that is capable of stating which formulas of the imperative language (corresponding to future constraints on planning) are in force at later states of the world, supposing the execution of some course of action from the present state of the world. It is also, of course, useful for our object language to be capable of representing, at least crudely, some kind of temporality—minimally, some sort of distinction between an agent's plans for the "present" (and the constraints impinging on her planning at present) and an agent's plans for different points in the future (and the constraints impinging on her planning at different points in the future). The language we define in this section is capable of all this, and more.

43. In designing the language and giving its semantics, we follow Segerberg (1990) closely, although there are some crucial differences. For example: (i) Segerberg allows imperative formulae to be embedded under the scope of Boolean operators, and we do not; (ii) Segerberg utilizes a monadic imperative operator, handling conditional imperatives in terms of Boolean $\rightarrow$, while we continue to insist on a dyadic imperative operator. Corresponding differences manifest in the semantics.

4.1.1 The Imperative Language $\mathcal{L}_{ILA}$

We first augment the alphabet, so that it consists of:

(87) **Alphabet.**
The Boolean propositional language $\mathcal{L}_P$.
A dyadic imperative operator $!(\cdot)(\cdot)$
An action operator δ
A modular modal operator $[\cdot]$
The regular operations $+$ and ;

We construct a set of *terms* $\mathcal{T}_{ILA}$ recursively.

(88) $\mathcal{T}_{ILA}$ is the smallest set such that:
If $\phi \in \mathcal{L}_P$, then $\delta\phi \in \mathcal{T}_{ILA}$.
If $\alpha, \beta \in \mathcal{T}_{ILA}$, then $\alpha;\beta \in \mathcal{T}_{ILA}$.
If $\alpha, \beta \in \mathcal{T}_{ILA}$, then $\alpha + \beta \in \mathcal{T}_{ILA}$.

Terms are the action expressions of the language. As before $\delta\phi$ is interpreted as *the action of seeing to it that* ϕ. Roughly, $\alpha;\beta$ designates the complex action of performing α then β, while $\alpha+\beta$ designates the complex action of performing at least one of α or β. The imperative language of action $\mathcal{L}_{ILA}$ is defined recursively as follows.

(89) **Definition of $\mathcal{L}_{ILA}$.**
If $\phi \in \mathcal{L}_P$, then $\phi \in \mathcal{L}_{ILA}$
If $\phi \in \mathcal{L}_P$ and $\alpha \in \mathcal{T}_{ILA}$, then $!(\phi)(\alpha) \in \mathcal{L}_{ILA}$
If $\alpha \in \mathcal{T}_{ILA}$ and $\phi \in \mathcal{L}_{ILA}$, then $[\alpha]\phi \in \mathcal{L}_{ILA}$.
Nothing else in $\mathcal{L}_{ILA}$.

Two comments about this new language. First, the left argument of the imperative operator continues to function as a restriction argument, while the right argument is now filled by an action, rather than a proposition. This is a natural amendment: imperatives command (and permit) *actions* in certain situations. Second, the language comes stocked with an infinite supply of dynamic modal operators $\{[\alpha] \mid \alpha \in \mathcal{T}_{ILA}\}$. The modal formula $[\alpha]\phi$ is to be read as *in all the states accessible by executing* α*,* ϕ. Note that dynamic modal operators can have imperative complements. This fact will be of considerable use to us in giving a semantics and pragmatics for complex imperatives of the form $!(\phi)(\alpha;\beta)$.

4.1.2 The Deontic Language $\mathcal{L}_{DLA}$

To build this language, we replace the two-place imperative operator $!(\cdot)(\cdot)$ with a two-place deontic necessity operator $O(\cdot)(\cdot)$ in the alphabet. The recursive definition of the language is unchanged, except for the replacement of the imperative clause with a corresponding clause for $O(\cdot)(\cdot)$, and a clause allowing Boolean combinations of arbitrary formulas of $\mathcal{L}_{DLA}$ (which we do not bother stating).

(90) If $\phi \in \mathcal{L}_{DLA}$ and $\alpha \in \mathcal{T}_{ILA}$, then $O(\phi)(\alpha) \in \mathcal{L}_{DLA}$.

Note that the prohibition on imperative formulas occurring in restrictor position is not extended to the deontic language: deontic formulas may function as restriction arguments for O.

This requires modifying the set of terms for $\mathcal{L}_{DLA}$—$\mathcal{T}_{DLA}$—specifically, it requires allowing actions to be made out of arbitrary formulas of the language.

(91) $\mathcal{T}_{DLA}$ is the smallest set such that:
If $\phi \in \mathcal{L}_{DLA}$, then $\delta\phi \in \mathcal{T}_{DLA}$.
If $\alpha, \beta \in \mathcal{T}_{DLA}$, then $\alpha;\beta \in \mathcal{T}_{DLA}$.
If $\alpha, \beta \in \mathcal{T}_{DLA}$, then $\alpha + \beta \in \mathcal{T}_{DLA}$.

Clearly, there is no one-one mapping from $\mathcal{T}_{ILA}$ to $\mathcal{T}_{DLA}$—expressions like $\delta!(\pi)(\delta\phi)$ are not terms of $\mathcal{T}_{ILA}$, while expressions like $\delta O(\pi)(\delta\phi)$ do count as terms of $\mathcal{T}_{DLA}$.

4.2 Models and To-Do Lists

We complicate models to accommodate the extra complexity of the languages.

(92) $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V \rangle$
D, W, V are typed as before.
$\mathcal{A} \subseteq 2^{W \times W}$ is a set of actions (relations on states).
$\Delta : 2^W \mapsto \mathcal{A}$ corresponds to the denotation of δ.
$\mathcal{C} \subseteq D \times D \times \mathcal{I} \times \mathcal{T}$, where $\mathcal{I} = \{i \mid i \subseteq W\}$, $\mathcal{T} = \{t \mid t \subseteq \mathcal{A}\}$

There are two changes worth noting. First, we add an algebra of actions for the sake of interpreting the practical part of our language (and a function mapping from propositions into this algebra). Actions are typed as relations on states. Conceiving of the elements of the universe as states of the world, rather than "whole" worlds, is conceptually perhaps significant, but mathematically insignificant. Doing so allows us to build a sort of temporality into the logic—actions are understood to relate prior states of the world to posterior states of the world, which will frequently differ in the formulas they satisfy. This enhances the logic's realism, while avoiding the complexities of temporal logic.

Second, we will be provisionally replacing the conversational backgrounds of the earlier models with simpler entities: world-invariant sets of worlds (rather than functions from worlds into sets of propositions) in the case of the modal base; world-invariant sets of actions, in the case of the ordering-source. For our purposes, letting the modal base be a world-invariant set of worlds, rather than a genuine Kratzer-ian conversational background is harmless. Letting the ordering-source be a world-invariant set of actions turns out to be less innocuous, but has a temporary dialectical justification—something close to it is a casual assumption in recent work on the pragmatics of imperatives (see, e.g., Isaacs & Potts 2003; Portner 2004, 2008; Potts 2003), and we will want to show why it needs to be discarded.[44]

Because our focus has shifted away from speaker commitments to planning (addressee constraints), we need to reconceive the context parameter of the semantics.

(93) A context $c \in \mathcal{C} = \langle s_c, a_c, i_c, t_c \rangle$
$i_c \in \mathcal{I}$ is the relevant information (modal base) in c
$t_c : D \mapsto \mathcal{I}$ is a function from individuals to their To-Do Lists in c

To-Do Lists are glossed in different ways in the contemporary literature on the pragmatics of imperatives. Ninan (2005) (following Portner 2004), for example, conceives them as sets of intentions (in our setup, sets of intended actions).[45] Portner (2008), on the

44. See §4.10 for further discussion of these points.

45. Or, more accurately, sets of public intentions—propositions such that it is Common Ground in c that an individual intends to fulfill them. This difference is unimportant for our purposes here.

other hand, glosses them as sets of actions to which an individual is committed.[46] Ninan's gloss falters on the fact that To-Do Lists have an *action-constraining* role to play in our modeling; it is not clear how an agent's actual intentions are supposed to constrain her planning, if at all (cf. the classic discussion of "bootstrapping" in Bratman 1987). Portner's gloss does better—an agent's commitments clearly generate constraints on her planning—although the intended sense of "commitment" is unclear and ripe for idiosyncratic interpretation. I propose a more minimal understanding of To-Do Lists for present purposes, along the lines of the "command sets" of Segerberg (1990). To-Do Lists just are sets of requirements on the agent's planning. They do not generate constraints; they just *are* constraints. Further development of the formal apparatus will clarify what is meant by this.

4.3 Semantics for Action Formulas

Satisfaction conditions for Boolean formulas are trivial and we omit them here. In giving the semantics for action formulas,[47] we begin by extending the interpretation function $[\![\cdot]\!]_{\mathcal{M},c}$ (mapping from pieces of the relevant language to their intensions, relative to a model $\mathcal{M}$ and context c) to cover terms of our languages. Let $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V\rangle$ be a model and c a context in $\mathcal{C}$. Then:

(94) a. $[\![\delta\phi]\!]_{\mathcal{M},c} = \Delta[\![\phi]\!]_{\mathcal{M},c}$
b. $[\![\alpha+\beta]\!]_{\mathcal{M},c} = [\![\alpha]\!]_{\mathcal{M},c} \cup [\![\beta]\!]_{\mathcal{M},c}$
c. $[\![\alpha;\beta]\!]_{\mathcal{M},c} = [\![\alpha]\!]_{\mathcal{M},c} \circ [\![\beta]\!]_{\mathcal{M},c}$

Note: $\Phi \circ \Psi = \{\langle w,v\rangle \mid \exists u : \langle w,u\rangle \in \Phi \wedge \langle u,v\rangle \in \Psi\}$. We lay down a statement of satisfaction conditions for the "action-representing" parts of $\mathcal{L}_{ILA}$ and $\mathcal{L}_{ILA}$ simultaneously, saving the semantics for imperative and deontic formulae until we have a bit more apparatus in place. Let $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V\rangle$ be a model, c a context in $\mathcal{C}$, and w a state in W. Then:

(95) $\mathcal{M}, c, w \Vdash_{\mathcal{L}_{ILA}} / \vDash_{\mathcal{L}_{DLA}} [\alpha]\phi$ iff
$\forall v : \langle w,v\rangle \in [\![\alpha]\!]_{\mathcal{M},c} \Rightarrow \mathcal{M}, c, v \Vdash_{\mathcal{L}_{ILA}} / \vDash_{\mathcal{L}_{DLA}} \phi$

We will want to enforce the following conditions on the Δ operation. Let $P \subseteq W$.

(96) a. If $\langle w,v\rangle \in \Delta P$, then $v \in P$.
[Equivalent to axiomatizing with $[\delta\phi]\phi$]
b. If, for any v, $\langle w,v\rangle \in \Delta P \Rightarrow v \in P'$,
then $\langle w,u\rangle \in \Delta P \Rightarrow \langle w,u\rangle \in \Delta P'$.
[Equivalent to axiomatizing with $[\delta\phi]\psi \rightarrow ([\delta\psi]\pi \rightarrow [\delta\phi]\pi)$]

As is easy to check, we have the following formulae valid in the class of all models for the logic.

(97) a. $[\alpha+\beta]\phi \leftrightarrow [\alpha]\phi \wedge [\beta]\phi$
[$\alpha+\beta$ always terminates in a ϕ-state iff both α and β do]
b. $[\alpha;\beta]\phi \leftrightarrow [\alpha][\beta]\phi$
[$\alpha;\beta$ always terminates in a ϕ-state iff every α-accessible state w is such that every state β-accessible from w is a ϕ-state]

46. Portner demurs from giving a realistic treatment of action, choosing to use properties as stand-ins for actions in his account. As we see, modeling certain phenomena about the imperative requires something a bit more subtle.

47. This part of the semantics is simply lifted from Segerberg (1990: 206-9).

4.4 Conditions on Models

Before giving the semantics for the "difficult" parts of our languages, we want to make sure that the models—their To-Do List parameters, in particular—are well-behaved. We will restrict the class of models for our languages to those satisfying the minimal conditions put forward in this section. Let $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V \rangle$ be a model, $c = \langle s_c, a_c, i_c, t_c \rangle$ be a context in $\mathcal{C}$, and d be an arbitrary individual in D. We first require that constraints never conflict with one another.

(98) If $\Phi_1 \in t_c(d)$ and $\Phi_2 \in t_c(d)$, then $\Phi_1 \cap \Phi_2 \neq \emptyset$ [**CON**]

While an agent can have incompatible desires, we suppose that planning constraints of the sorts enforced by imperatives must be consistent with one another.[48]

Additionally, we require that any non-absurd action on a To-Do List at a context be available—possible, that is to say, given the context-relevant information. There are two ways to enforce the intuition behind this requirement.

(99) If $\Phi \in t_c(d)$ and $\Phi \neq \emptyset$, then, for some $\langle w, v \rangle \in \Phi$, $w \in i_c$ [**AV1**]

(100) For some $w \in i_c$, v: if $\Phi \in t_c(d)$ and $\Phi \neq \emptyset$, then, $\langle w, v \rangle \in \Phi$ [**AV2**]

We temporarily regard information states as supplying the relevant possible input states for an agent's action. (As we see in §4.7, it is somewhat more perspicuous to regard them as supplying information about more than possible input states.) **AV1** has it that an action cannot occur on a To-Do List if the information rules out the possibility of it being performed. For example, an agent cannot bound by a constraint to *maintain* ϕ unless some input states are ϕ-states. **AV2** is a strengthening of **AV1**: it has it that every action on the To-Do List is *simultaneously executable* at some state compatible with the c-relevant information. **AV2**, then, amounts to a global realism constraint on To-Do Lists (or, alternatively, an ideality constraint on modal bases): constraints must be simultaneously satisfiable with respect to the information at the context. **AV2** is plausible. Supposing for now that constraints on To-Do Lists reliably generate corresponding obligations—i.e., true descriptions of obligation—and that obligations aggregate, this is a natural assumption to make. So we will make it.[49]

We might contemplate adding one of the following two closure conditions:

(101) If $\Phi_1 \in t_c(d)$, $\Phi_2 \in t_c(d)$, and $\Phi_1 \cap \Phi_2 \subseteq \Phi_3$, then $\Phi_3 \in t_c(d)$ [**CL1**]

(102) If $\Phi_1 \in t_c(d)$, $\Phi_2 \in t_c(d)$, then $\Phi_1 \cap \Phi_2 \in t_c(d)$ [**CL2**]

CL1 requires, inter alia, that To-Do Lists are closed under intersection, union, and arbitrary expansion (so it is stronger than **CL2**). **CL1** entails, inter alia, that constraints on planning both distribute and aggregate (as, e.g., in 103a and 103b), while **CL2** entails that they aggregate (as, e.g., in 103b).

(103) a. If $[\![\delta(\phi \rightarrow \psi)]\!]_{\mathcal{M},c} \in t_c(d)$ and $[\![\delta\phi]\!]_{\mathcal{M},c} \in t_c(d)$, then $[\![\delta\psi]\!]_{\mathcal{M},c} \in t_c(d)$
b. If $[\![\delta\phi]\!]_{\mathcal{M},c} \in t_c(d)$ and $[\![\delta\psi]\!]_{\mathcal{M},c} \in t_c(d)$, then $[\![\delta(\phi \wedge \psi)]\!]_{\mathcal{M},c} \in t_c(d)$

48. Because, as we shall soon see, To-Do Lists are in the business of directly generating obligations for an agent, **CON** amounts to a To-Do List analogue of the **D** axiom: $O\phi \rightarrow \neg O \neg \phi$.

49. As we will be using To-Do Lists as ordering-sources for deontic formulas of our new deontic object language, **AV2** is also a strengthened version of the Limit Assumption. Instead of merely requiring that there be best worlds in the modal base, **AV2** will require that there be ideal worlds—worlds compatible with the execution of every constraint—in the modal base. Again, while this would not be plausible for a bouletic ordering-source (in view of incompatible desires), it is plausible for an ordering-source constituted by binding constraints.

It is natural to have the following set of intuitions about these conditions. Thinking ahead, and exploiting the sorts of intuitions that Segerberg (1990) uses to motivate his above-discussed restrictions on Command Systems, we might object to **CL1**, on the grounds that, down the line, there is reason to expect it to cause difficulty for attempting a resolution the Ross Paradox. We might also endorse **CL2**, on the grounds that the constraint on planning that imperatives enforce intuitively ought to aggregate. These intuitions presuppose that To-Do Lists have no role to play in a planning semantics for imperatives, besides being devices for keeping track of requirements on agents. The presupposition is, as we are about to see, inapt. And as we see a bit later on: (i) our preferred resolution of the Ross Paradox in this framework cares naught about closure conditions on To-Do Lists (§4.11); (ii) **CL2** is unnecessary to secure the desired results about aggregation of obligations and constraints on planning. **CON** and **AV2** are sufficient to get a semantics for the deontic and imperative action languages going. Simplicity, when available, is a virtue, and we will stick with them for the moment.

We begin building such a semantics now. Note that the analyses presented in the subsequent few sections are only a first pass, and modifications will have to be made to accommodate conditional imperatives, the Ross Paradox, and "higher-order" types of imperative. As in our discussion in the prior half of the paper, we will introduce these modifications piecemeal, as the occasion demands.

4.5 Orderings on Transitions

In shifting our interest from a logic of content to a logic of planning, we also implicitly shifted the sort of deontic modality we would be interested in modeling. For most of this paper, we have been interested in a kind of bouletic deontic modality—obligation in the view of the desires of a potential issuer of an imperative. In formulating a logic of planning, it is natural to shift our attention to obligation in view of the constraints that bear on an agent's planning.

Going with what has worked for us, this naturally means allowing To-Do Lists to fill the semantic role we have allowed speaker-related bouletic ordering-sources to play.[50] Since To-Do Lists are typed differently than ordering-sources—the former being sets of relations on W, the latter sets of subsets of W—the appropriate extension is not immediate.

We first use To-Do Lists to define a preorder on $W \times W$. The resulting ordering is on inter-state transitions, rather than worlds. Let c be an arbitrary context, and d an arbitrary individual. (Implicitly, these definitions are all with respect to a model, but we will continue to economize on notation by suppressing this.)

(104) $\langle v, v' \rangle \leq_{t_c(d)} \langle u, u' \rangle$ iff
$\{\mathbf{\Phi} \in t_c(d) \mid \langle u, u' \rangle \in \mathbf{\Phi}\} \subseteq \{\mathbf{\Phi} \in t_c(d) \mid \langle v, v' \rangle \in \mathbf{\Phi}\}$

A transition $\langle v, v' \rangle$ is at least as good as a transition $\langle u, u' \rangle$ with respect to a To-Do List $t_c(d)$ iff for every action $\mathbf{\Phi}$ in $t_c(d)$ of which $\langle u, u' \rangle$ is a transition, $\langle v, v' \rangle$ is a transition of $\mathbf{\Phi}$.

In a logic of action, it is most natural to treat modals as quantifiers over inter-state transitions, rather than states simpliciter. In this vein, we define the domain of the deontic necessity modal with respect to a set of input states (information-state) p, set of output

50. Portner (2008) is, to my knowledge, the first (i) to notice the natural connection between To-Do Lists and the interpretation of certain commitment-describing root modals, and (ii) to propose accounting for this connection by utilizing To-Do Lists as the ordering-sources for such modals.

states q, and To-Do List $t_c(d)$—$\mathit{min}[p \times q, \leq_{t_c(d)}]$—as follows.

(105) $\mathit{min}[p \times q, \leq_{t_c(d)}]$ is the set of all $\langle v, v' \rangle \in p \times q$ such that:
$\forall \langle u, u' \rangle \in p \times q : \langle u, u' \rangle \leq_{t_c(d)} \langle v, v' \rangle \Rightarrow \langle v, v' \rangle \leq_{t_c(d)} \langle u, u' \rangle$

The best state-to-state transitions $\langle v, v' \rangle$ with respect to a set of input states p, output states q, and To-Do List $t_c(d)$ are those such that (i) their input state is a relevant possible input state (i.e., $v \in p$); (ii) their output state is a possible output state (i.e., $v' \in q$); (iii) for any other transition $\langle u, u' \rangle$ satisfying (i) and (ii), either $\langle u, u' \rangle$ is strictly worse (with respect to $t_c(d)$) than $\langle v, v' \rangle$, or they are equally good.

We give two passes at stating satisfaction conditions for deontic formulas of $\mathcal{L}_{DLA}$ utilizing this setup: one that explicitly restricts the set of transitions from which best transitions are drawn to nomologically admissible transitions (equivalently: explicitly restricts the set of possible output states, relative to a body of law), and one which does not.

4.6 Explicitly Restricted

We explicitly define a notion of possible output states, relative to a set of possible input states. Let $\mathit{nom}_{L,d} \subseteq W \times W$ be the set of transitions that are nomologically admissible for an agent d, according to some body of law L (perhaps encoding information about the laws of nature, causal limitations of the agent, etc.). Roughly, if $\langle v, v' \rangle \in \mathit{nom}_{L,d}$, then the agent *can* effect a transition to v', provided she finds herself in v.[51] We define the set of *relevant possible actions* for d—the set of actions whose input states are possible relative to an information state p and the relevant body of law—as follows:

(106) $\mathit{rel}(p,d) = (p \times W) \cap \mathit{nom}_{L,d}$

Our first pass at stating satisfaction conditions for deontic formulae of $\mathcal{L}_{DLA}$ is as follows:

(107) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ iff
$\mathit{min}[\mathit{rel}(i_c \cap [\![\phi]\!]_{\mathcal{M},c}, a_c), \leq_{t_c(a_c)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$

$O(\phi)(\alpha)$ is satisfied in a model $\mathcal{M}$ at a context-world pair $\langle c, w \rangle$ iff all the best (with respect to a_c's To-Do List) transitions **from** ϕ-states in i_c **to** possible output states (states to which a_c can effect transitions from the ϕ-states in i_c) are transitions of α.

An illustration will be useful. Let $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V \rangle$ be a model, $c \in \mathcal{C}$, and let:

- $W = \{w_1, ..., w_6\}$
- $[\![\phi]\!]_{\mathcal{M},c} = \{w_1, w_3, w_5\}$ (ϕ-states indicated by circled nodes)
- $i_c = \{w_1, w_2, w_3, w_5\}$ (indicated with straight arrows)
- $t_c(a_c)$ be the closure of $\{\{\langle w_2, w_4 \rangle, \langle w_3, w_4 \rangle, \langle w_5, w_4 \rangle, \langle w_6, w_4 \rangle\}\}$ under the conditions on To-Do Lists specified above (squiggly arrows diagram $\bigcap t_c(a_c)$)

51. It would be good to say something precise about the character of the body of law and how the body of law restricts the set of available actions for an agent, but doing so is *very* difficult.

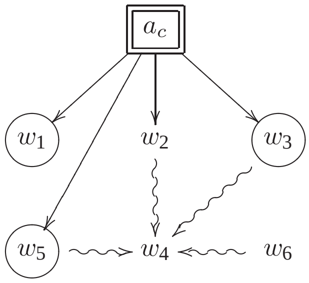

Let $[\![\alpha]\!]_{\mathcal{M},c} = \{\langle w_1, w_4\rangle, \langle w_2, w_4\rangle, \langle w_3, w_4\rangle\}$. To evaluate the truth of $O(\phi)(\alpha)$ in $\mathcal{M}$ at $\langle c, w\rangle$ (for arbitrary w), we fix our attention on the ϕ-states in i_c.

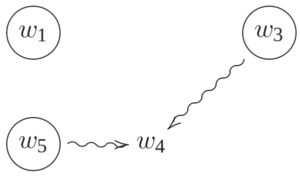

Suppose w_4 is nomologically accessible from each of w_1, w_3, w_5 (but they are not accessible from each other), so that *rel*$(i_c \cap [\![\phi]\!]_{\mathcal{M},c}, a_c) = \{\langle w_1, w_4\rangle, \langle w_3, w_4\rangle, \langle w_5, w_4\rangle\}$. Then $O(\phi)(\alpha)$ is true in $\mathcal{M}$ at $\langle c, w\rangle$ iff *min*$[\{\langle w_1, w_4\rangle, \langle w_3, w_4\rangle, \langle w_5, w_4\rangle\}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$. But the best transitions are precisely those labeled on the last graph (recall, squiggly arrows diagram $\bigcap t_c(a_c)$): $\langle w_3, w_4\rangle$ and $\langle w_5, w_4\rangle$. Since $\langle w_5, w_4\rangle \notin [\![\alpha]\!]_{\mathcal{M},c}$, $O(\phi)(\alpha)$ is false in $\mathcal{M}$ at $\langle c, w\rangle$.

4.7 Implicitly Restricted

The nomological accessibility semantics is interesting and potentially useful, which is why we have taken time to formulate it: it provides a way to integrate, in an explicit way, information about the causal structure of the world into the semantics for statements of obligation. Given that we have nothing to say about the relevant body of law, however, it is a complication to no immediate benefit. The semantics for deontic formulae we develop henceforth will be an extension of the semantics given in this section.

Since we are ultimately understanding deontic operators as universal quantifiers over transitions, it is profitable to reconceive the modal base i_c as a set of transitions, rather than states. Let $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V\rangle$ be a model. We re-type the $\mathcal{C}$ parameter and the modal base, as follows:

(108) $\mathcal{C} \subseteq D \times D \times \mathcal{I} \times \mathcal{T}$,
where $\mathcal{I} = \{i \mid i \subseteq W \times W\}$

(109) A context $c \in \mathcal{C} = \langle s_c, a_c, i_c, t_c\rangle$
$i_c \in \mathcal{I}$ is the relevant information (modal base) in c

Information states are thus reconceived as encoding information both about the possible current states of the world and the causal structure of the world. If $\langle v, v'\rangle \in i_c$, then, for all that is known by the lights of the c-relevant information, the current state of the world is v, and v' is a possible successor-state of v. Information states are treated as embodying information about the current state of the world and what sorts of changes to the world are possible to effect. This change requires a trivial reformulation of the **AV** constraints on models. The new constraints will motivated in the same way as **AV**. Regarding **AV1***:

actions must be possible with respect to the relevant information. And regarding **AV2***: actions must be simultaneously executable with respect to the relevant information.

(110) If $\Phi \in t_c(d)$ and $\Phi \neq \emptyset$, then, for some $\langle w,v \rangle \in \Phi$, $\langle w,v \rangle \in i_c$ [**AV1***]

(111) For some $\langle w,v \rangle \in i_c$: if $\Phi \in t_c(d)$ and $\Phi \neq \emptyset$, then, $\langle w,v \rangle \in \Phi$ [**AV2***]

We give our second pass at stating satisfaction conditions for deontic formulae of $\mathcal{L}_{DLA}$:

(112) $\mathcal{M},c,w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ iff
$\mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$

$O(\phi)(\alpha)$ is satisfied in a model $\mathcal{M}$ at a context-world pair $\langle c,w \rangle$ iff all the best (with respect to a_c's To-Do List) transitions in i_c having ϕ-states as inputs are transitions of α.

4.8 Semantics for Imperatives

Now that we have something like a fully-formed semantics on the table, we may formulate requirement conditions for imperative formulas of $\mathcal{L}_{ILA}$. We continue our practice of giving requirement conditions for imperative formulas in terms of satisfaction conditions for deontic formulas. Let $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V \rangle$ be a model, c a context in $\mathcal{C}$, and w a state in W. Then:

(113) $\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} !(\phi)(\alpha)$ iff $\mathcal{M},c,w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$

4.9 Further Conditions on Models

We may obtain interesting sub-logics by further constraining models. Some worth considering:

(114) If $\Phi \in t_c(a_c)$, then $\Phi \subseteq i_c$ [**NT**]

(115) If $\Phi \in t_c(a_c)$, then $i_c \subseteq \Phi$ [**KE**]

(116) If $\Phi \in t_c(a_c)$, then $\{w' \mid \exists w : \langle w,w' \rangle \in i_c\} \subseteq \{w' \mid \exists w : \langle w,w' \rangle \in \Phi\}$ [**KE1**]

(117) If $\Phi \in t_c(a_c)$, then $\{w \mid \exists w' : \langle w,w' \rangle \in i_c\} \subseteq \{w \mid \exists w' : \langle w,w' \rangle \in \Phi\}$ [**KE2**]

Each has an air of plausibility, but only one is a reasonable addition to our standing set of constraints. Since we are working with a rather non-standard apparatus, it will be useful to talk through all of them. We discuss them in sequence.

4.9.1 Non-Triviality

Condition **NT** is a non-triviality requirement on transitions of actions on To-Do Lists: they cannot suggest inter-state transitions whose possibility is ruled out by the c-relevant information. Non-triviality requirements have a degree of plausibility about them, especially in the context of a logic of planning: a To-Do List should not be able to suggest a transition that the agent is unable, by the lights of the available information, to execute. Nevertheless, we should tread carefully. First, **NT** has questionable motivation. Although we might loosely speak of agent's executing transitions between states, it is important to realize that what an agent actually executes is an action. The transition is the side-effect of the execution of an action, just as the transition to a new memory state is the side-effect of the execution of computer's execution of a program. Insofar as the non-triviality intuition exerts any real pull, then, it concerns actions: actions must be executable by the light of the information available at the context; there are no trivial (non-executable) programs on an agent's To-Do List. But **AV1*** and **AV2*** already succeed in accommodating this incarnation of the intuition. Less abstractly, we know that in the dynamic (diachronic) context, information acquisition regularly *shrinks* the set of transitions compatible with the relevant information, while adding nothing new to the constraints impinging on the agent's planning (although it does frequently impact an agent's planning about how to satisfy *prior* constraints). Modeling how information acquisition impacts planning is difficult without an account on which agents might somehow satisfy a non-triviality requirement "on the fly." Without some reconception of certain parameters of the semantics, information acquisition will regularly bring an agent into violation of **NT**. This will require either an adjustment in her To-Do List (but why should the set of constraints that *structure* planning be altered by information acquisition?) or an addition of states to i_c (but why should information acquisition regularly be accompanied with inexplicable information leakage?).

4.9.2 Knowledge of Executability

Conditions **KE**, **KE1**, and **KE2** are reminiscent of the condition on ordering-sources stated in our previous discussion of monotonicity (§3.4.5). Each is an attempts to formalize the intuition that actions on To-Do Lists must be known to be executable, by the lights of the c-relevant information. According to **KE**, no action $\Phi \in t_c(a_c)$ is such that there is a transition compatible with the c-relevant information that is not a transition of Φ. According to **KE1**, no action $\Phi \in t_c(a_c)$ is such that there is an output state compatible with the c-relevant information that is not an output state of Φ. According to **KE2**, no action $\Phi \in t_c(a_c)$ is such that there is an input state compatible with the c-relevant information that is not an input state of Φ.

KE turns out to be untenable in our system: it implies that $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ iff $i_c \cap (\llbracket\phi\rrbracket_{\mathcal{M},c} \times W) \subseteq \llbracket\alpha\rrbracket_{\mathcal{M},c}$. That is to say: α is obligatory (supposing ϕ) iff α is knowably executable (supposing ϕ).

> Proof. $\Leftarrow$ is immediate. So suppose $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$.
>
> Then $min[i_c \cap (\llbracket\phi\rrbracket_{\mathcal{M},c} \times W), \leq_{t_c(a_c)}] \subseteq \llbracket\alpha\rrbracket_{\mathcal{M},c}$.
>
> By **KE**, $i_c \subseteq \bigcap t_c(a_c)$. So $\bigcap t_c(a_c) \cap i_c = i_c$.
>
> Then, by def. of *min*, $min[i_c \cap (\llbracket\phi\rrbracket_{\mathcal{M},c} \times W), \leq_{t_c(a_c)}] = i_c \cap (\llbracket\phi\rrbracket_{\mathcal{M},c} \times W)$.
>
> So $i_c \cap (\llbracket\phi\rrbracket_{\mathcal{M},c} \times W) \subseteq \llbracket\alpha\rrbracket_{\mathcal{M},c}$.

Informally, then, **KE** has it that any action that is possible for the agent to take, by the lights of the available information, conditional on whatever, will satisfy all of the constraints binding on her at c. This evidently trivializes the action-guiding role that To-Do Lists are conceived to have in planning. Similarly, **KE1** has it that any informationally possible action will *fulfill* all of the constraints binding on the agent at c; there is no informationally possible way for the agent to go wrong.

KE2 does a fine job enforcing knowable executability of actions on To-Do Lists, provided we understand the modal base in the correct way. **KE2** has it, roughly, that any possible way for the *current* state of the world to be is a way does not rule out the performance of any action on the agent's To-Do List. This means understanding the modal base as having a dual role: as representing *epistemic* uncertainty about the current state of the world and *nomological* uncertainty or possibility about the future state of the world (rather than epistemic uncertainty about nomological possibility): if $\langle w, v\rangle \in i_c$, then for all the information available at c, w might be the current state of the world, and v is nomologically accessible (given the relevant laws) from w. On this understanding of the modal base, **KE2** together with **AV1*** suffice to secure knowable executability in the relevant sense. **KE2** has it that the agent is, by the lights of the information about the current state of the world, in a position to execute any action on her To-Do List. **AV1*** has it that, for any action on an agent's To-Do list, it is nomologically possible for her to fulfill that action's purpose; whatever the current state w of the world, should she choose to execute any action on her To-Do List in w, she can expect that execution to terminate in a new state of the world in which that action is fulfilled.

It is important to note that **KE2** cannot play the same role in a PDLA as our restriction of ordering-sources for deontic modals to sets of *choices* (knowably actionable desires). Recall that the latter restriction was used to secure a non-monotonicity property for the set of best worlds. If all actions in the ordering-source-like To-Do List are antecedently required to be knowably executable, shrinking the modal base i_c will not alter the membership of the To-Do List. Let $i_{c'} \subseteq i_c$ be an arbitrary restriction of the modal base.

> Proof. Suppose that **KE2** and that $\Phi \in t_c(a_c)$.
> Then $\{w \mid \exists w' : \langle w, w' \rangle \in i_c\} \subseteq \{w \mid \exists w' : \langle w, w' \rangle \in \Phi\}$.
> Since $i_{c'} \subseteq i_c$, $\{w \mid \exists w' : \langle w, w' \rangle \in i_{c'}\} \subseteq \{w \mid \exists w' : \langle w, w' \rangle \in i_c\} \subseteq \{w \mid \exists w' : \langle w, w' \rangle \in \Phi\}$.

Augmenting the context with information does not, by the lights of **KE2**, affect whether an item on a To-Do List is knowably executable. So **KE2** *cannot be used to secure non-monotonicity*. Nevertheless, the assumption of **KE** makes non-monotonicity superfluous in the semantics. Unless the action described by *save all ten miners* is knowably executable, in the sense specified by **KE2**, it cannot appear on a To-Do List, and cannot serve to privilege transitions that terminate in its execution. (In case the action *is* knowably executable, then we should want claims like *you ought to either block A or block B* to come out true regardless.) While this sort of treatment is implausible for speaker-given bouletic ordering-sources (it is possible to desire something without knowing how to fulfill your desire), it seems acceptable for To-Do lists. If you are genuinely constrained to perform some action, it ought to be known that the world will allow you to execute it, should you attempt to.[52]

4.10 Contingency

While it may seem like we have a working PDLA semantics for both qualified statements of obligation (in view of imperative constraints) and conditional imperatives, the unfortunate truth is that we have nothing of the sort. The function of a conditional imperative $(\textit{if } \phi)(\textit{stit } \psi)$ is to embody constraints on an agent's planning in input-states satisfying ϕ. Similarly, the function of qualified statement of constraint-governed obligation $(\textit{if } \phi)(\textit{ought } \psi)$ is to make a statement about what constraints require in input-states satisfying ϕ. But according to the semantics we have given for conditional imperatives and

52. In spite of my breezy tone, the territory is *very* thorny, and the discussion here cannot do it justice. It is, in fact, important for the semantics to have access to something like non-monotonicity in the case of conditional commands, although for something like the opposite of the reasons of Kolodny & MacFarlane (2008). Namely (and rather astonishingly): restricting i_c to $i_{c'}$ can make actions that were knowably executable with respect to i_c *fail to be knowably executable* with respect to $i_{c'}$. Pursuing this topic would unfortunately involve a complication of the semantics that would muddy the waters later on. Things are going to get fairly technical—the simpler the apparatus, the easier it will be to follow what's going on. So we bracket it, with the hope of pursuing it in further research.

conditional statements of constraint-governed obligation, the function of these constructions is something else entirely—a conditional imperative embodies a categorical, or non-contingent, constraint on an agent's planning behavior, while a conditional statement of constraint-governed obligation makes a claim about an agent's categorical obligations, her obligations come what may.

Recall the statement of satisfaction conditions in item (112) for an arbitrary deontic formula in the deontic language of action $O(\phi)(\alpha)$, repeated here.

(118) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ iff
$\mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$

Note the *non-occurrence of the index variable* w in the statement of satisfaction-conditions: the values of the relevant parameters—the modal base i_c and ordering-source $t_c(a_c)$—are both fixed at the context c, and cannot be shifted by shifting the index of evaluation. Supposing that $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$, it follows immediately that $\mathcal{M}, c, v \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$, <u>for all $v \in W$</u>. There are no contingently true statements of obligation—statements of the form $O(\phi)(\alpha)$ such that they are satisfied in some worlds (with respect to a model and a context) and not in others. Because of the equivalence between satisfaction-conditions for deontic formulas of $\mathcal{L}_{DLA}$ and requirement-conditions for imperative formulas of $\mathcal{L}_{ILA}$, this infects the semantics for imperatives, in a slightly altered fashion: commands are always categorically required (i.e., with respect to every point of evaluation).

I take it as a datum there are contingently true statements of obligation (and contingently required commands), and, further, that this is precisely what conditional statements of obligation (and conditional imperatives) are in the business of expressing. The formula $O(\phi)(\alpha)$ expresses that the best ϕ-initial transitions are transitions of α, i.e., that the execution of α is required should one find oneself in a ϕ state from which such a transition is executable, i.e., that $O(\top)(\alpha)$ is satisfied in every ϕ state from which such a transition is executable, although not necessarily at the original point of evaluation. This intuition is, I think, especially grabbing for the logician of planning. If every ϕ-initial transition τ where the agent meets all of the constraints on her action is a transition of α, then the agent may reasonably take herself as required to perform α, if she finds herself in a position to execute τ.

The model theory we have formulated is inadequate to this task, and we will devote some amount of effort to revising it. The basic problem is building index-dependence into the To-Do List.[53] But this is not a simple matter of construing To-Do Lists as Kratzerian conversational backgrounds, so that $t_c(a_c) : W \mapsto 2^{\mathcal{A}}$ (although we will end up doing this). While doing so makes the semantics *compatible* with contingent obligation and contingent requirement, it will not reliably *predict* that $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ implies $\mathcal{M}, c, v \vDash_{\mathcal{L}_{DLA}} O(\top)(\alpha)$, for all v such that $\langle v, v' \rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}]$.

4.10.1 Building in Index-Sensitivity

We begin by re-typing To-Do Lists roughly as conversational backgrounds, so that $t_c(a_c) : W \mapsto 2^{\mathcal{A}}$. Informally, To-Do Lists will be reconceived as functions from worlds to sets of constraints binding at those worlds ("contingent constraints" for short).[54] We begin at ground level, with models. With this machinery on the table, we will churn quickly

53. While the modal base ought to be made index-dependent, that does not help with the problem here. Let in_{i_c} designate $\{w \mid \exists w' : \langle w, w' \rangle \in i_c\}$ (the set of inputs to i_c). A well-behaved modal base—one obeying Reflexivity ($w \in \mathsf{in}_{i_c}(w)$) and Euclideanness (if $v \in \mathsf{in}_{i_c}(w)$, then $\mathsf{in}_{i_c}(w) \subseteq \mathsf{in}_{i_c}(v)$) constraints—is, for practical purposes, index-invariant. Together, Reflexivity and Euclideanness entail: if $v \in \mathsf{in}_{i_c}(w)$, $\mathsf{in}_{i_c}(w) = \mathsf{in}_{i_c}(v)$. See Gillies (2008) for a proof.

54. To-Do Lists are thus typed in the same way as the Command Systems of Segerberg (1990).

through our "axiom" system, figuring out how to rewrite our constraints on models to accommodate the new apparatus. Formally, let $\mathcal{M} = \langle D, W, \mathcal{A}, \Delta, \mathcal{C}, V \rangle$ be a model, with all parameters except $\mathcal{C}$ (and all parameters of $\mathcal{C}$ save $\mathcal{T}$) typed as in §4.7.

(119) $\mathcal{C} \subseteq D \times D \times \mathcal{I} \times \mathcal{T}$,
where $\mathcal{T} = \{t \mid t : D \mapsto \{l \mid l : W \mapsto 2^{\mathcal{A}}\}\}$

(120) A context $c \in \mathcal{C} = \langle s_c, a_c, i_c, t_c \rangle$

Requirements on models are restated as follows, with informal glosses appended.

(121) If $\Phi_1 \in t_c(d)(w)$ and $\Phi_2 \in t_c(d)(w)$, then $\Phi_1 \cap \Phi_2 \neq \emptyset$ [**CON***]
[If Φ_1 and Φ_2 are constraints at w, Φ_1 and Φ_2 are compatible.]

(122) For some $\langle w, v \rangle \in i_c$: if $\Phi \in t_c(d)(w)$ and $\Phi \neq \emptyset$, then, $\langle w, v \rangle \in \Phi$ [**AV2****]
[Constraints at w are simultaneously executable, by lights of i_c.]

(123) If $\Phi \in t_c(a_c)(w)$, then $\{v \mid \exists v' : \langle v, v' \rangle \in i_c\} \subseteq \{v \mid \exists v' : \langle v, v' \rangle \in \Phi\}$ [**KE2***]
[Contingent constraints are knowably executable, by lights of i_c.]

Concerning **AV2****, it may be implausible to think that contingent constraints that hold at states that are incompatible with i_c's picture of the current state of the world should always be realizable, by the lights of i_c. Since (on account of reflexivity of the modal base—see fn53) we will only be interested in evaluating deontic formulas at states of evaluation that are compatible with the modal base's picture of the current state of the world, we can avoid the implausibility (and lose nothing of any importance) by holding that $t_c(d)(w)$ is defined just in case $w \in \{w \mid \exists w' : \langle w, w' \rangle \in i_c\}$. This gets around the problem, such as it is.

There are two natural ways to utilize our new To-Do Lists in giving a satisfactory semantics for deontic formulas of $\mathcal{L}_{DLA}$: treating them as ordering-sources (and giving an ordering-source semantics) and, equivalently (given certain natural constraints on To-Do Lists), treating them neighborhood functions (and giving a neighborhood semantics). I will sketch each way in brief.

4.10.2 Sensitive Ordering-Source Semantics

Defining an index-relative preorder on transitions and set of good-enough transitions is only a matter of extending prior definitions, and I will spare the reader repetition. We set forth new satisfaction conditions for deontic formulas—first, the general case, second, the special case where the restriction argument of O is vacuous.

(124) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ iff
$\mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$

(125) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\top)(\alpha)$ iff
$\mathit{min}[i_c, \leq_{t_c(a_c)(w)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$

The change to satisfaction-conditions is minimal—most of the work having gone into making the ordering-source index-sensitive. Note that the $\mathcal{L}_{DLA}$ version of **K** is $\mathcal{L}_{DLA}$-valid (by the lights of the ordering-source semantics) in the class of all models for the language.

(126) $O(\pi)(\delta(\phi \to \psi)) \to (O(\pi)(\delta\phi) \to O(\pi)(\delta\psi))$ [**K**]

Proof. Suppose $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\pi)(\delta(\phi \to \psi))$ and $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\pi)(\delta\phi)$. Then each best-at-$w$ π-initial transition is a transition of both $\delta(\phi \to \psi)$ and $\delta\phi$, and thereby of $\delta\psi$.

While the semantics is, as it stands, compatible with contingent obligation (and contingent requirement), we have work to do to secure the desired prediction: $\mathcal{M},c,w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ should imply $\mathcal{M},c,v \vDash_{\mathcal{L}_{DLA}} O(\top)(\alpha)$, for all v such that $\langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}]$. But this can be achieved by adding a minimal further constraint on To-Do Lists—one which would not have been possible on the original treatment of To-Do Lists as index-insensitive. (We save discussion of this constraint for §4.10.3.)

(127) If $\mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(d)(w)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$, then if $\langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(d)(w)}]$, $[\![\alpha]\!]_{\mathcal{M},c} \in t_c(d)(v)$.

Informally, if the best transitions in a set of transitions i, relative to $t_c(d)(w)$, are transitions of α, then α is a contingent constraint, in force at every input-state of every best transition.[55] It is fairly easy to show that (127) secures the desired result.

Proof. We suppose that $\mathcal{M},c,w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$.

Then $\mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(d)(w)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$.

Then, if $\langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(d)(w)}]$, $[\![\alpha]\!]_{\mathcal{M},c} \in t_c(d)(v)$.

Suppose $\langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(d)(w)}]$. Then $[\![\alpha]\!]_{\mathcal{M},c} \in t_c(d)(v)$.

By **AV2****, $\exists \langle u,u'\rangle \in i_c : \langle u,u'\rangle \in \Phi$, for each $\Phi \in t_c(d)(v)$.

Then $\forall \langle t,t'\rangle \in \mathit{min}[i_c, \leq_{t_c(d)(v)}] : \langle t,t'\rangle \in \Phi$ for each $\Phi \in t_c(d)(v)$.

So $\langle t,t'\rangle \in [\![\alpha]\!]_{\mathcal{M},c}$, for any $\langle t,t'\rangle \in \mathit{min}[i_c, \leq_{t_c(d)(v)}]$.

So, $\mathcal{M},c,v \vDash_{\mathcal{L}_{DLA}} O(\top)(\alpha)$.

4.10.3 Sensitive Neighborhood Semantics

Alternatively, we can state an index-sensitive neighborhood semantics for deontic formulas of $\mathcal{L}_{DLA}$, utilizing To-Do Lists more or less as neighborhood functions.

(128) $\mathcal{M},c,w \vDash_{\mathcal{L}_{DLA}} O(\phi)(\alpha)$ iff
$\forall \langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] : [\![\alpha]\!]_{\mathcal{M},c} \in t_c(a_c)(v)$

(129) $\mathcal{M},c,w \vDash_{\mathcal{L}_{DLA}} O(\top)(\alpha)$ iff
$\forall \langle v,v'\rangle \in \mathit{min}[i_c, \leq_{t_c(a_c)(w)}] : [\![\alpha]\!]_{\mathcal{M},c} \in t_c(a_c)(v)$ iff
$\forall \langle v,v'\rangle \in \bigcap t_c(a_c)(w) \cap i_c : [\![\alpha]\!]_{\mathcal{M},c} \in t_c(a_c)(v)$

Informally, $O(\phi)(\alpha)$ is satisfied in $\mathcal{M}$ at $\langle c,w\rangle$ iff α is a contingent constraint holding at all the best ϕ-initial transitions. As with our earlier pass at a neighborhood semantics (§3.5.1), the To-Do List still plays a role in selecting best transitions, but, additionally, directly tells us what sorts of constraints are in force at each state.

All that is needed to establish equivalence between the neighborhood and ordering-source formulations of the semantics is the addition of a further minimal constraint on To-Do lists.

(130) If $\forall \langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(d)(w)}]$, $[\![\alpha]\!]_{\mathcal{M},c} \in t_c(d)(v)$, then $\mathit{min}[i_c \cap ([\![\phi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(d)(w)}] \subseteq [\![\alpha]\!]_{\mathcal{M},c}$.

Note that this condition is simply the converse of (127). Taken together, (127) and (130) amount to to endorsing a limited *indifference condition* on obligations: an agent has a conditional obligation at w just in case the unconditional obligation is in force at every

55. There is no way to express this constraint in the object language, for reasons I will not go into here.

point v such that (i) v satisfies the relevant condition and (ii) v allows execution of each of the constraints binding on her at w. To be a bit impressionistic, on this picture, we have it that obligations are both *contingent* and *coarse*. Contingent because they may, of course, vary between points of evaluation. Coarse because having a conditional obligation at w means having the unconditional obligation at all states that satisfy the relevant condition and do not vary in the actions an agent can, in view of the information at her disposal, execute from them.

4.11 Incorporating Permission

In §3.5.1, we argued against neighborhood semantic resolutions of the Ross Paradox, on the grounds that their failure to explicitly represent permissive content led to explanatory gaps. We have proposed a sort of neighborhood semantics for $\mathcal{L}_{DLA}$ (and indirectly $\mathcal{L}_{ILA}$), but not on the grounds that it provides any sort of satisfactory resolution of the Ross Paradox (indeed, we will see presently that it does not). The virtue of neighborhood semantic treatments—and the source of my interest in them—is their intuitiveness and theoretical simplicity, especially in the case of constraint-describing deontic formulas and constraint-expressing imperative formulas, as well as the nice fit the neighborhood semantics will have with imperative dynamics (see §5.3). A deontic formula $O(\phi)(\alpha)$ (and imperative formula $!(\phi)(\alpha)$) is satisfied (required) just in case α *appears on the agent's To-Do List*, in every ϕ-state where the agent is in a position to meet every applicable constraint.

The semantics, as it stands, renders $O(\pi)(\delta\phi) \rightarrow O(\phi)(\delta\phi + \delta\psi)$ valid in every model.

> Proof. Suppose that $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\pi)(\delta\phi)$.
> Then $\mathit{min}[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] \subseteq [\![\delta\phi]\!]_{\mathcal{M},c}$.
> Note that $[\![\delta\phi]\!]_{\mathcal{M},c} \subseteq [\![\delta\phi + \delta\psi]\!]_{\mathcal{M},c}$.
> So, $\mathit{min}[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] \subseteq [\![\delta\phi + \delta\psi]\!]_{\mathcal{M},c}$
> So, $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\pi)(\delta\phi + \delta\psi)$.

It thus predicts that $!(\pi)(\delta\phi) \Vdash_{\mathcal{L}_{ILA}} !(\pi)(\delta\phi + \delta\psi)$. That is to say, it predicts the Ross Paradoxical inference—its $\mathcal{L}_{ILA}$ representation, anyway—valid.

But that is fine, so far as it goes. The apparatus we have developed is flexible, and is, suitably modified, capable of making the reverse prediction—doing so, moreover, by having imperatives play both constraining and *permitting* roles in planning behaviors of agents. We will suggest two ways of modifying the semantics here. The first is in roughly the same vein as the alternative semantics for disjunctive imperatives developed in §3.5.2. The second makes use of alternatives in representing the permissive role of imperatives, but introduces a novel formal setup—namely, a complication of the context parameter with a device for explicitly represents the *rights* or practical *entitlements* of individuals—to integrate them into the semantics.

4.11.1 Alternative Semantics

We might well treat "disjunctive" action-terms as inducing alternatives, along the lines of the account developed in §3.5.2. We would define salient alternatives for imperative formulae of $\mathcal{L}_{ILA}$, doing so recursively as follows.

(131) a. $\mathit{alt}[!(\pi)(\delta\phi)] = \delta\phi$
b. $\mathit{alt}[!(\pi)(\alpha_1; ...; \alpha_n)] = \mathit{alt}[!(\pi)(\alpha_1)]; ...; \mathit{alt}[!(\pi)(\alpha_n)]$
c. $\mathit{alt}[!(\pi)(\alpha_1 + ... + \alpha_n)] = \mathit{alt}[!(\pi)(\alpha_1)], ..., \mathit{alt}[!(\pi)(\alpha_n)]$

A few comments about this.

i. We let action-terms function as salient alternatives, rather than formulas.

ii. We assume that a *no-choice* interpretation of a disjunctive imperative $(\textit{if}\ \pi)(\textit{stit}(\phi \vee \psi))$ will be assigned the $\mathcal{L}_{ILA}$ logical form $!(\pi)(\delta(\phi \vee \psi))$. *Free choice* interpretations are mapped to the $\mathcal{L}_{ILA}$ logical form $!(\pi)(\delta\phi + \delta\psi)$.

    This is intuitive: no-choice interpretations have $(\textit{if}\ \pi)(\textit{stit}(\phi \vee \psi))$ suggesting a single means of compliance if π (i.e., $\delta(\phi \vee \psi)$), while free-choice interpretations have it suggesting two means of compliance if π (i.e., both $\delta\phi$ and $\delta\psi$). Since we use distinct logical forms to handle the relevant interpretations, rather than allowing context to decide an extra-logical set of salient alternatives, we no longer require that salient alternatives be relativized to models, contexts, and worlds. (It is, I think, a nice side effect of our adopting a formal imperative language capable of representing actions that we can do things this way.[56])

iii. Finally, we interpret $!(\pi)(\alpha_1;...;\alpha_n)$ as inducing an *ordered sequence* of alternatives $\textit{alt}[!(\pi)(\alpha_1)];...;\textit{alt}[!(\pi)(\alpha_n)]$. What this amounts to will emerge in a bit. Only note that it does *not* amount to introducing ordered-sequences into either of our object languages.

We proceed to define relevant permission operators $[\wp]$, $[P]$, and P in the deontic language. We introduce an abbreviation to make notation less cumbersome: if $\boldsymbol{A}$ is a list of actions $\beta_1,...,\beta_n$, then let $\bigoplus \boldsymbol{A} = \beta_1 + ... + \beta_n$. To appropriately define $[\wp]$, we require some way of talking about which formulas *hold true* in the best worlds meeting some condition or other (which we had lost upon moving to our new languages). We therefore augment the object languages with a modular restrictable dynamic modal operator $[\![\cdot]\!]^{\cdot}(\cdot)$, such that if α is an action term of the relevant language, π is a non-imperative formula of the relevant object language, and ϕ is any formula of the relevant object language, then $[\![\alpha]\!]^{\pi}(\phi)$ is a formula of the relevant object language. We give this modal operator the following semantics.

(132) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} / \Vdash_{\mathcal{L}_{ILA}} [\![\alpha]\!]^{\pi}(\phi)$ iff
$\forall \langle v, v' \rangle \in \textit{min}[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W) \cap [\![\alpha]\!]_{\mathcal{M},c}, \leq_{t_c(a_c)(w)}] : \mathcal{M}, c, v' \vDash_{\mathcal{L}_{DLA}} / \Vdash_{\mathcal{L}_{ILA}} \phi$

$[\![\alpha]\!]^{\pi}(\phi)$ says the best π-initial executions of α yield a state satisfying (requiring) ϕ. Some facts worth noting about this new operator. Clearly, the relevant instance of **K** is valid in the class of all $\mathcal{L}_{DLA}$ models.

(133) $[\![\alpha]\!]^{\pi}(\phi \rightarrow \psi) \rightarrow ([\![\alpha]\!]^{\pi}(\phi) \rightarrow [\![\alpha]\!]^{\pi}(\psi))$ **[K#]**

We note that the following formula is also valid in the class of all $\mathcal{L}_{DLA}$ models. Note: we use $\mho$ to denote the vacuous action, so that $[\![\mho]\!]_{\mathcal{M},c} = W \times W$.

(134) $[\![\mho]\!]^{\pi}(\phi) \leftrightarrow O(\pi)(\delta\phi)$

Proof. $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} [\![\mho]\!]^{\pi}(\phi)$ iff
$\forall \langle v, v' \rangle \in \textit{min}[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] : \mathcal{M}, c, v' \vDash_{\mathcal{L}_{DLA}} \phi$ iff

56. To note, then set aside: our objection of §3.5.2 to Aloni (2007)'s approach to the Ross Paradox—that it is eliminativist about imperative logic and attributes the Ross Paradox to undesirable coarseness in the imperative object language—does not apply to the account we give here. Obviously, our account is non-eliminativist. Further, $\mathcal{L}_{ILA}$ is fine-grained enough to represent the differences between free- and no-choice interpretations of disjunctive imperatives. $\mathcal{L}_{DLK}$ is not, but that is not the object language we are using at the moment.

$min[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] \subseteq [\![\delta\phi]\!]_{\mathcal{M},c}$ iff
$\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\pi)(\phi)$.

Note that this result does not hold generally, i.e., we do not have that, for all substitutions of action terms for α, $[\![\alpha]\!]^{\pi}(\phi) \leftrightarrow O(\pi)(\delta\phi)$.

We use our new operator to define the critical permission operator $[\wp]$.

(135) Permissions in $\mathcal{L}_{DLA}$.
a. $[\wp][\pi][\boldsymbol{A}_1 \ ; \ ... \ ; \boldsymbol{A}_n] :=$
$[P](\pi)(\boldsymbol{A}_1) \wedge ... \wedge [\![\bigoplus \boldsymbol{A}_1 \ ; \ ... \ ; \bigoplus \boldsymbol{A}_{n-1}]\!]^{\pi}([P](\top)(\boldsymbol{A}_n))$
b. $[P](\pi)(\alpha_1, ..., \alpha_n) := P(\pi)(\alpha_1) \wedge ... \wedge P(\pi)(\alpha_n)$
c. $P(\pi)(\alpha) := \neg O(\pi)(\overline{\alpha})$

$[\wp][\pi][\boldsymbol{A}_1; ...; \boldsymbol{A}_n]$ says, roughly, that any action in $\boldsymbol{A}_1$ is okay if π; and any ideal π-initial execution of any action in $\boldsymbol{A}_1$ yields a state in which any action in $\boldsymbol{A}_2$ is okay; and ... ; and, finally, any ideal π-initial execution of any action in $\boldsymbol{A}_1$, then any action in $\boldsymbol{A}_2$, then ..., then, finally, any action in $\boldsymbol{A}_{n-1}$ yields a state in which any action in $\boldsymbol{A}_n$ is okay. Informally, $[\wp][\pi][\boldsymbol{A}_1; ...; \boldsymbol{A}_n]$ expresses trickling down of permissions through an ordered sequence of alternatives, provided that initial permissions are executed in accordance with planning constraints.

To illustrate, consider the complex imperative $!(\pi)(\alpha_1 + \alpha_2; \beta_1 + \beta_2)$ (read: if π, do either α_1 or α_2, then do either β_1 or β_2). Our definitions yield the following.

- $alt[!(\pi)(\alpha_1 + \alpha_2; \beta_1 + \beta_2)] := \alpha_1, \alpha_2 \ ; \ \beta_1, \beta_2$.
- $[\wp](\pi)(\alpha_1, \alpha_2 \ ; \ \beta_1, \beta_2) := [P](\pi)(\alpha_1, \alpha_2) \wedge [\![\alpha_1 + \alpha_2]\!]^{\pi}([P](\top)(\beta_1, \beta_2)) :=$
$P(\pi)(\alpha_1) \wedge P(\pi)(\alpha_2) \wedge [\![\alpha_1 + \alpha_2]\!]^{\pi}(P(\top)(\beta_1) \wedge P(\top)(\beta_2))$

Informally, $[\wp](\pi)(\alpha_1, \alpha_2 \ ; \ \beta_1, \beta_2)$ expresses that α_1 and α_2 are permitted if π, and that, in all the situations that can be accessed via a good enough (by the lights of the To-Do List), π-initial execution of α_1 or α_2, β_1 and β_2 are permitted (simpliciter). This seems like the right result. Here we are beginning to see some real benefits to adopting languages of action—we can model permissions of arbitrarily complex, temporal imperative constructions, and track how permissions trickle down to subsequent situations, contingent on an agent's behavior in antecedent situations.

To connect this to the Ross Paradox, we need to connect salient alternatives to requirement conditions for imperatives. We do this by redefining the ∇ operator, and revise requirement conditions for imperatives, as follows.

(136) $\nabla(\pi)(\boldsymbol{A}_1 \ ; \ ... \ ; \boldsymbol{A}_n) :=$
$[\wp][\pi][\boldsymbol{A}_1 \ ; \ ... \ ; \boldsymbol{A}_n] \wedge O(\pi)(\bigoplus \boldsymbol{A}_1; ...; \bigoplus \boldsymbol{A}_n)$

(137) If $\phi = !(\pi)(\alpha)$:
$\mathcal{M}, c, w \Vdash_{\mathcal{L}_{ILA}} \phi$ iff $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} \nabla(\pi)(alt(\phi))$

Clearly, this setup makes the right predictions about the most basic version of the Ross Paradox: we easily predict that $!(\pi)(\delta\phi) \nVdash_{\mathcal{L}_{ILA}} !(\pi)(\delta\phi + \delta\psi)$. $!(\pi)(\delta\phi + \delta\psi)$ is required at $\mathcal{M}$ in $\langle c, w \rangle$ only when we have $P(\pi)(\delta\psi)$ satisfied at $\mathcal{M}$ in $\langle c, w \rangle$. Clearly, though, for some $\mathcal{M}$ and $\langle c, w \rangle$, $\mathcal{M}, c, w \Vdash !(\pi)(\delta\phi)$ and $\mathcal{M}, c, w \nvDash_{\mathcal{L}_{DLA}} P(\pi)(\delta\psi)$. And so we have it that $!(\pi)(\delta\phi) \nVdash_{\mathcal{L}_{ILA}} !(\pi)(\delta\phi + \delta\psi)$.

This solution is readily extendable to arbitrarily complicated versions of the Ross Paradox. As an illustration: we automatically predict that $!(\pi)(\alpha_1 + \alpha_2; \beta_1) \nVdash_{\mathcal{L}_{ILA}} !(\pi)(\alpha_1 +$

$\alpha_2;\beta_1+\beta_2)$. The latter imperative expresses a permission to perform β_2, supposing admissible (by the lights of the To-Do List) execution of $\alpha_1+\alpha_2$ in a π-initial state. The former imperative does not. So $!(\pi)(\alpha_1+\alpha_2;\beta_1) \not\Vdash_{\mathcal{L}_{ILA}} !(\pi)(\alpha_1+\alpha_2;\beta_1+\beta_2)$.

Note that if we had chosen to endorse **LC***, we would now have a contradiction on our hands. Suppose $\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} !(\pi)(\delta\phi)$. Then, by the ordering-source semantics stated in §4.10.2, we have that $\forall\langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] : [\![\delta\phi]\!]_{\mathcal{M},c} \in t_c(a_c)(v)$. Let $\langle v,v'\rangle \in \mathit{min}[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}]$. **LC*** has To-Do Lists closed under arbitrary expansion, so that $[\![\delta\phi]\!]_{\mathcal{M},c} \cup [\![\delta\psi]\!]_{\mathcal{M},c} \in t_c(a_c)(v)$. So $[\![\delta\phi+\delta\psi]\!]_{\mathcal{M},c} \in t_c(a_c)(v)$. So, by the neighborhood semantics stated in §4.10.3, $\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} !(\pi)(\delta\phi+\delta\psi)$. But then $!(\pi)(\delta\phi) \Vdash_{\mathcal{L}_{ILA}} !(\pi)(\delta\phi+\delta\psi)$ after all!

The benefits of this approach speak for themselves. Nevertheless, there is a downside. While we have a device in the formal apparatus for keeping track of constraints introduced on an agent's planning behavior over time (the To-Do List), we have no such device for keeping track of *rights* or *entitlements*. Such a device is inessential for handling the Ross Paradox, but rather useful for modeling the changes in an agent's planning behavior introduced by a commanding authority: authorities impose constraints, but also grant freedoms. We will introduce it now, show that it can handle the Ross Paradox—indeed, give a resolution equivalent to the one just developed—but save an account of its role in the pragmatics of imperatives for later.

4.11.2 Explicitly Representing Rights

The relevant modification of the apparatus is simple and intuitive. We introduce a "dual" parameter to an agent's To-Do List: a Rights List. The function of a Rights List is to keep track of the actions an agent is entitled to perform at different states of the world. The notion of a right or entitlement is related to, but, significantly, *not exhausted by*, the notion of an action an agent is not constrained not to perform. An agent has a right to perform an action, roughly, just in case an authority has granted her such a right. While every right of an agent will correspond to an action the agent is not constrained not to perform, not every action the agent is not constrained not to perform will correspond to a right of an agent. Let $\mathcal{M} = \langle D,W,\mathcal{A},\Delta,\mathcal{C},V\rangle$ be a model. We re-type the $\mathcal{C}$ parameter in the expected way.

(138) $\mathcal{C} \subseteq D \times D \times \mathcal{I} \times \mathcal{T} \times \mathcal{T}$

(139) A context $c \in \mathcal{C} = \langle s_c, a_c, i_c, t_c, r_c\rangle$
$r_c : D \mapsto \{l \mid l : W \mapsto 2^{\mathcal{A}}\}$ is a Rights List function for c

Rights, naturally, are allowed to vary according to index in the same way as constraints. Intuitively, $r_c(d)(w)$ will yield a set of actions—*contingent rights*—for an agent d at state w.

Derivative Rights Lists. Rights Lists are built up from an entity we will refer to as a Derivative Rights List (DRL). A DRL r_c^- is characterized in terms of To-Do lists (whence their "derivative" moniker).

(140) $r_c^-(d)(w) =_{df} \{\Phi \mid \overline{\Phi} \notin t_c(d)(w)\}$

DRLs, then, give the set of actions an agent is not constrained not to perform at an index of evaluation. A property of DRLs worth noting:

(141) If $\Phi_1 \in r_c^-(d)(w)$, then $\Phi_1 \cap \Phi_2 = \emptyset$ implies $\Phi_2 \notin t_c(d)(w)$

Suppose for reductio that $\Phi_1 \in r_c^-(d)(w)$, $\Phi_1 \cap \Phi_2 = \emptyset$ and $\Phi_2 \in t_c(d)(w)$. Because To-Do Lists are closed under arbitrary expansion, and because $\Phi_2 \subseteq \overline{\Phi_1}$, we have $\overline{\Phi_1} \in t_c(d)(w)$. But then, by (140), $\Phi_1 \notin r_c^-(d)(w)$. Contradiction. Informally, this means that actions appearing on DRLs are always compatible with each of the constraints binding on an agent at w. And because To-Do lists are closed under intersection, this in turn implies that every action on an agent's DRL at w is compatible with simultaneous fulfillment of all of the constraints binding on an agent at w.[57]

This setup allows us to state a neighborhood semantics in terms of DRLs for permission formulas of $\mathcal{L}_{DLA}$. Recall the neighborhood semantics of §4.10.3, which yields a statement of satisfaction conditions for permission formulas in terms of To-Do Lists:

(142) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} P(\phi)(\alpha)$ iff
$\exists \langle v, v' \rangle \in \mathit{min}[i_c \cap (\llbracket \phi \rrbracket_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] : \llbracket \overline{\alpha} \rrbracket_{\mathcal{M},c} \notin t_c(a_c)(v)$

Given (140)'s definition of the DRL, the following conditions are equivalent.

(143) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} P(\phi)(\alpha)$ iff
$\exists \langle v, v' \rangle \in \mathit{min}[i_c \cap (\llbracket \phi \rrbracket_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] : \llbracket \alpha \rrbracket_{\mathcal{M},c} \in r_c^-(a_c)(v)$

The resolution of the Ross Paradox proposed in the prior section can be rendered in this setup, by giving a neighborhood semantics for $[\wp][\pi][\boldsymbol{A}_1 \,;\, \ldots \,;\, \boldsymbol{A}_n]$ in terms of DRLs:

(144) $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} [\wp][\pi][\boldsymbol{A}_1 \,;\, \ldots \,;\, \boldsymbol{A}_n]$ iff

- $\alpha \in \boldsymbol{A}_1 \Rightarrow \exists \langle v, v' \rangle \in \mathit{min}[i_c \cap (\llbracket \pi \rrbracket_{\mathcal{M},c} \times W), \leq_{t_c(a_c)(w)}] : \llbracket \alpha \rrbracket_{\mathcal{M},c} \in r_c^-(a_c)(v)$ &
- $\alpha \in \boldsymbol{A}_2 \Rightarrow \forall \langle v, v' \rangle \in \mathit{min}[i_c \cap (\llbracket \pi \rrbracket_{\mathcal{M},c} \times W) \cap \llbracket \bigoplus \boldsymbol{A}_1 \rrbracket_{\mathcal{M},c}, \leq_{t_c(a_c)(w)}] :$
  $\exists \langle u, u' \rangle \in \mathit{min}[i_c, \leq_{t_c(a_c)(v')}] : \llbracket \alpha \rrbracket_{\mathcal{M},c} \in r_c^-(a_c)(u)$ &

  $\vdots$

- $\alpha \in \boldsymbol{A}_n \Rightarrow \forall \langle v, v' \rangle \in \mathit{min}[i_c \cap (\llbracket \pi \rrbracket_{\mathcal{M},c} \times W) \cap \llbracket \bigoplus \boldsymbol{A}_1; \ldots; \bigoplus \boldsymbol{A}_{n-1} \rrbracket_{\mathcal{M},c}, \leq_{t_c(a_c)(w)}] :$
  $\exists \langle u, u' \rangle \in \mathit{min}[i_c, \leq_{t_c(a_c)(v')}] : \llbracket \alpha \rrbracket_{\mathcal{M},c} \in r_c^-(a_c)(u)$

Non-derivative Rights Lists. We introduce Non-derivative Rights Lists (NDRLs) here, saving application and elucidation for later on (§5.4). NDRLs are conceived as restrictions of DRLs:

(145) If $\Phi \in r_c(d)(w)$, then $\Phi \in r_c^-(d)(w)$.

We introduce several constraints on NDRLs.

(146) If $\Phi \in t_c(d)(w)$, then $\Phi \in r_c(d)(w)$ **[CR]**

(147) If $\Phi_1 \in r_c(d)(w)$, then $\Phi_1 \cap \Phi_2 = \emptyset$ implies $\mathit{min}[i_c, \leq_{t_c(d)(w)}] \not\subseteq \Phi_2$ **[NON]**

CR has it that contingent constraints are also contingent rights—items on the To-Do list count as rights of the agent. **NON** (mnemonic for *not obligated not*) has it that an agent's rights at w are not contradicted by contrary obligations at w (cf. fn57).

57. This is *not* to say that if $\llbracket \alpha \rrbracket_{\mathcal{M},c} \in r_c^-(d)(w)$, then $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} P(\top)(\alpha)$. (This is, then, the source of a disanalogy between DRLs and To-Do lists.) If $\llbracket \alpha \rrbracket_{\mathcal{M},c} \in r_c^-(d)(w)$, it is the case that $\bigcap t_c(d)(w) \not\subseteq \llbracket \overline{\alpha} \rrbracket_{\mathcal{M},c}$. But it still may be the case that $\bigcap t_c(d)(w) \cap i_c \subseteq \llbracket \overline{\alpha} \rrbracket_{\mathcal{M},c}$. Perhaps the only ideal transitions, by the lights of i_c, are transitions of $\llbracket \overline{\alpha} \rrbracket_{\mathcal{M},c}$. We could get around this by letting $r_c^-(d)(w) =_{df} \{\Phi \cap i_c \mid \overline{\Phi} \notin t_c(d)(w)\}$. This would require that, if $\llbracket \alpha \rrbracket_{\mathcal{M},c} \in r_c^-(d)(w)$, then $\bigcap t_c(d)(w) \cap i_c \cap \llbracket \alpha \rrbracket_{\mathcal{M},c} \neq \emptyset$.

NDRLs are constrained in some of the same ways as To-Do Lists, with some important differences. We will not generally expect them to obey Rights List versions of **LC***, **CON***, or **AV2***. (The same, incidentally, goes for DRLs.)

(148) If $\Phi_1 \in r_c(d)(w)$, $\Phi_2 \in r_c(d)(w)$, and $\Phi_1 \cap \Phi_2 \subseteq \Phi_3$, then $\Phi_3 \in r_c(d)(w)$ [**LC-R**]

(149) If $\Phi_1 \in r_c(d)(w)$ and $\Phi_2 \in r_c(d)(w)$, then $\Phi_1 \cap \Phi_2 \neq \emptyset$ [**CON-R**]

(150) For some $\langle v, v' \rangle \in i_c$: if $\Phi \in t_c(d)(w)$ and $\Phi \neq \emptyset$, then, $\langle v, v' \rangle \in \Phi$ [**AV2-R**]

Concerning **LC-R**, an agent can have, for example, a right to post the letter and a right to not post the letter, but no agent has a single right to both post and not post the letter. Similar considerations tell against requiring Rights Lists to meet constraints **CON-R** and **AV2-R**: agents may naturally have freedoms (though not obligations) that are impossible to simultaneously execute.

Analogues of the **AV**** and **KE*** constraints are more reasonable. We will also require that NDRLs be closed under arbitrary expansion of actions (**EX-R**).

(151) If $\Phi \in r_c(d)(w)$ and $\Phi \neq \emptyset$, then, for some $\langle v, v' \rangle \in \Phi$, $\langle v, v' \rangle \in i_c$ [**AV1-R**]

(152) If $\Phi \in r_c(d)(w)$, then $\{v \mid \exists v' : \langle v, v' \rangle \in i_c\} \subseteq \{v \mid \exists v' : \langle v, v' \rangle \in \Phi\}$ [**KE2-R**]

(153) If $\Phi_1 \in r_c(d)(w)$, then if $\Phi_1 \subseteq \Phi_2$, $\Phi_2 \in r_c(d)(w)$ [**EX-R**]

AV1-R requires that an agent has a right to α only if α is possibly executable, while **KE2-R** requires that rights be knowably executable. **AV-R** is a natural constraint to impose, given the relationship we will seek to codify between rights and $\mathcal{L}_{DLA}$ statements of permission: permissibility implies possibility with respect to the modal base. **KE-R** is less intuitive, but still rather natural: an agent has freedom to perform α only if α is knowably executable. Adopting **KE-R** makes **NON** superfluous (cf. again fn57). **EX-R** is a desirable closure constraint on NDRLs and allows us to state a satisfactory version of the dynamic entailment relation (§5.5).

Unlike DRLs, NDRLs do not have a role to play in the static semantics for permission formulas of our language. While every right is a permission, every permission is not necessarily a right. Their use is essentially dynamic—they will be used to keep track of rights introduced by authorities, in order to constrain the course of dynamic update. We return to this subject in §5.4.

4.12 Temporal Phenomena

As we've hinted, logics of action really come into their own with temporal imperative constructions. In this section, we explore several kinds of temporal phenomena about the imperative and show that the analysis we have been developing is well-suited to handling them.[58]

58. The problem of accounting for temporal imperative phenomena obviously also arises for the logician of content. While our motivation for adopting a logic capable of expressing facts about actions and their impact on the world was to account for the constraints on planning behavior enforced by imperatives, this fact suggests that we might use a logic of action to represent speaker commitments. I make a gesture at unified logic of content and planning in §4.13.

4.12.1 Ordered Commands

Instructions and commands are often order-sensitive. Your soufflé might well fail to rise if you happen to mistake the instruction in (154a) for the instruction in (154b).

(154) a. Whip the egg whites, then fold them into the custard. [*stit* ϕ *then* ψ]
b. Whip and fold the egg whites into the custard. [*stit*$(\phi \wedge \psi)$]

The logic we have been developing is, if course, well-suited to representing the difference between (154a) and (154b), assigning them the following $\mathcal{L}_{ILA}$ representations, respectively.

(155) a. $!(\top)(\delta\phi;\delta\psi)$
b. $!(\top)(\delta(\phi \wedge \psi))$ [or, perhaps, $!(\top)(\overline{\delta\neg\phi + \delta\neg\psi})$]

Since $\Delta([\![\phi]\!]_{\mathcal{M},c}) \circ \Delta([\![\psi]\!]_{\mathcal{M},c}) \nsubseteq \Delta([\![\phi \wedge \psi]\!]_{\mathcal{M},c})$ and $\Delta([\![\phi \wedge \psi]\!]_{\mathcal{M},c}) \nsubseteq \Delta([\![\phi]\!]_{\mathcal{M},c}) \circ \Delta([\![\psi]\!]_{\mathcal{M},c})$, there is no semantic relationship to speak of between these formulas of $\mathcal{L}_{ILA}$. This is a desirable feature for an imperative logic to exhibit, and it is all but automatic in the PDLA setup we have been using. So far as I can tell, the only way for a logic to mimic this result, without going in for the representation of action in the object language, is to make the object language explicitly temporal, rather than merely implicitly temporal—to introduce devices for talking directly about precedence, quantification over times (as we will see below), and the like. It is a virtue of the PDLA approach that we can represent rather complicated temporal phenomena with a rather minimal object language.

4.12.2 Temporal Constraints on Models

Because we have not yet laid down either object language axioms for ordered commands or corresponding conditions on To-Do Lists, our treatment of ordered commands is, as it stands, incomplete. We remedy this deficiency now, endorsing the following set of conditions on To-Do Lists. Let Ψ be an any restriction of i_c. Then:[59]

(156) If $min[\Psi, \leq_{t_c(d)(w)}] \subseteq \Phi_1 \circ \Phi_2$, then $min[\Psi, \leq_{t_c(d)(w)}] \subseteq \Phi_1$
[Equivalent to axiomatizing with $O(\phi)(\alpha;\beta) \rightarrow O(\phi)(\alpha)$]

(157) If $min[\Psi, \leq_{t_c(d)(w)}] \subseteq \Phi_1 \circ \Phi_2$ and $\langle v, v' \rangle \in min[\Psi \cap \Phi_1, \leq_{t_c(d)(w)}]$, then $min[i_c, \leq_{t_c(d)(v')}] \subseteq \Phi_2$
[Equivalent to axiomatizing with $O(\phi)(\alpha;\beta) \rightarrow [\![\alpha]\!]^{\phi}(O(\top)(\beta))$]

(158) If $min[\Psi, \leq_{t_c(d)(w)}] \subseteq \Phi_1$ and $\langle v, v' \rangle \in min[\Psi \cap \Phi_1, \leq_{t_c(d)(w)}] \Rightarrow min[i_c, \leq_{t_c(d)(v')}] \subseteq \Phi_2$, then $min[\Psi, \leq_{t_c(d)(w)}] \subseteq \Phi_1 \circ \Phi_2$
[Equivalent to axiomatizing with $(O(\phi)(\alpha) \wedge [\![\alpha]\!]^{\phi}(O(\top)(\beta))) \rightarrow O(\phi)(\alpha;\beta)$]

Endorsing condition (156) means we predict $!(\phi)(\alpha;\beta) \Vdash_{\mathcal{L}_{ILA}} !(\phi)(\alpha)$—correctly it seems, from the standpoint of a logic of planning. The constraint on planning enforced by $!(\phi)(\alpha;\beta)$ at a state w requires the performance of α, then β. So it requires, inter alia, the performance of α at w. Endorsing condition (157) means predicting that $!(\phi)(\alpha;\beta) \Vdash_{\mathcal{L}_{ILA}} [\![\alpha]\!]^{\phi}(!(\top)(\beta))$. If the performance of α, then β, is required at w if ϕ, then in the best ϕ-initial states where α is performed, β is required. This too is intuitive. Finally, endorsing (158) means we predict that $!(\phi)(\alpha), [\![\alpha]\!]^{\phi}(!(\top)(\beta)) \Vdash_{\mathcal{L}_{ILA}} !(\phi)(\alpha;\beta)$. If α is required, if ϕ, and all the best ϕ-initial executions of α yield states where β is required, then effect, α

59. These conditions on To-Do Lists are adaptations of conditions on Command Systems given by Segerberg (1990: 210).

then β is required if ϕ.[60]

Ordered commands, in conjunction with this set of reasonable temporal constraints on models, have interesting consequences for axiomatizing the logic, one of which we note here. Consider the following axiom for the dynamic operator $[\![\cdot]\!]^{\cdot}(\cdot)$.

(159) $\quad [\![\mho]\!]^{\pi}([\![\mho]\!]^{\pi}(\phi) \rightarrow \phi)$ **[TR]**

TR says that in all the best π-initial transitions: if all the best π-initial transitions satisfy ϕ, then ϕ. As such, it reads like an apparent $\mathcal{L}_{DLA}$ analogue of the familiar deontic axiom **OU**: $O(O\phi \rightarrow \phi)$ (it ought to be that, if ϕ is required, then ϕ). Not only does **TR** read like **OU**: it yields similar consequences together with the **K#** axiom for $[\![\cdot]\!]^{\cdot}(\cdot)$, repeated here.

(160) $\quad [\![\alpha]\!]^{\pi}(\phi \rightarrow \psi) \rightarrow ([\![\alpha]\!]^{\pi}(\phi) \rightarrow [\![\alpha]\!]^{\pi}(\psi))$ **[K#]**

OU together with the **K** axiom for the standard deontic language $O(\phi \rightarrow \psi) \rightarrow (O\phi \rightarrow O\psi)$ implies $OO\phi \rightarrow O\phi$. Similarly, note that the validity of **TR**, together with **K#**, immediately implies the validity of **CD4**.

(161) $\quad [\![\mho]\!]^{\pi}([\![\mho]\!]^{\pi}(\phi)) \rightarrow [\![\mho]\!]^{\pi}(\phi)$ **[CD4]**

Nevertheless, while **OU** is a reasonable axiom for the standard deontic language, the possibility of ordered commands tells decisively against axiomatizing our semantics with **TR**: **TR** implies that $O(\top)(\mho;\beta) \rightarrow O(\top)(\beta)$ is a universal validity—and thus that $!(\top)(\mho;\beta) \Vdash_{\mathcal{L}_{ILA}} !(\top)(\beta)$.

> Proof. Suppose **TR** is valid and $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\top)(\mho;\delta\psi)$.
> Then $min[i_c, \leq_{t_c(a_c)(w)}] \subseteq [\![\mho]\!]_{\mathcal{M},c} \circ [\![\delta\psi]\!]_{\mathcal{M},c}$.
> Let $\langle v, v'\rangle \in min[i_c, \leq_{t_c(a_c)(w)}]$. By (157), $min[i_c, \leq_{t_c(a_c)(v')}] \subseteq [\![\delta\psi]\!]_{\mathcal{M},c}$.
> Then $\mathcal{M}, c, v' \vDash_{\mathcal{L}_{DLA}} O(\top)(\delta\psi)$. By (134), $\mathcal{M}, c, v' \vDash_{\mathcal{L}_{DLA}} [\![\mho]\!]^{\top}(\psi)$.
> But then $\forall\langle v, v'\rangle \in min[i_c, \leq_{t_c(a_c)(w)}] : \mathcal{M}, c, v' \vDash_{\mathcal{L}_{DLA}} [\![\mho]\!]^{\top}(\psi)$.
> So $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} [\![\mho]\!]^{\top}([\![\mho]\!]^{\top}(\psi))$.
> Since **TR** is valid, **CD4** is too.
> So it follows that $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} [\![\mho]\!]^{\top}(\psi)$.
> Finally, by (134), $\mathcal{M}, c, w \vDash_{\mathcal{L}_{DLA}} O(\top)(\delta\psi)$.

This is backwards. An agent whose planning behavior is constrained by the imperative $!(\top)(\mho;\beta)$ should *not* necessarily look to perform β presently: the imperative prescribes performance of β only upon occasion of the agent's prior fulfillment of her present obligations.[61]

4.12.3 Stable Commands

There is a felt difference between what I will term *ephemeral commands*—commands whose influence on an agent's planning lapses upon fulfillment—and *stable commands*—

60. These results hold in the more complicated permission semantics for imperatives outlined in §4.11. We will simplify our discussions by suppressing some of the apparatus when it has no direct bearing on the matter at hand.

61. I note that we are also perhaps able to justice to Castaneda's intuitions about argument (17). If we represent the argument as having the form $[\alpha](!(\top)(\beta)), !(\top)(\alpha)\ /\ !(\top)(\beta)$, we see that the argument indeed comes out invalid, although not quite for the reasons that Castaneda happened to articulate.

commands whose influence on planning persists through instances of compliance.

(162) a. Buy groceries on Sunday! [*stit* ϕ]
b. Always buy groceries on Sunday! [*always stit* ϕ]

Intuitively, an ephemeral command *stit* ϕ is in force at w just in case all the To-Do List-preferred-at-w transitions terminate in a ϕ-state. By contrast, a stable command *always stit* ϕ is in force at w just in case all the To-Do List-preferred-at-w transitions terminate in a ϕ-state w' such that all the To-Do List-preferred-at-w' transitions terminate in a ϕ-state w'' such that... Our semantics, suitably adapted, is well-equipped to handle this intuitive semantic difference. Ephemeral commands *stit* ϕ are associated with the $\mathcal{L}_{ILA}$ formula $!(\top)(\delta\phi)$. That is to say, ephemeral commands are handled as garden-variety imperative constructions in our imperative object language; no additional machinery is required for their analysis.

Stable commands are trickier, but ultimately tractable, given a suitable extension of the machinery. It is natural, in view of the intuitive meaning of stable commands, to associate the stable command *always stit* ϕ with an infinitary formula $!(\top)((\delta\phi)^\star)$ (where $\star$ is the Kleene Star).[62] Alternatively, choosing to hew to finitary object languages, it would be natural to analyze *always stit* ϕ in terms of $!(\top)(\delta\phi^\omega)$, where ω is a finitary operation on action-terms, characterized as follows.

(163) If $\alpha \in \mathcal{T}_{ILA/DLA}$, then $\alpha^\omega \in \mathcal{T}_{ILA/DLA}$.

(164) $[\![\alpha^\omega]\!]_{\mathcal{M},c} = [\![\alpha^\star]\!]_{\mathcal{M},c} = ([\![\alpha]\!]_{\mathcal{M},c})^*$, where R^* is the ancestral of R.

Equivalently, we may analyze *always stit* ϕ in terms of dynamic modal operators: *always stit* ϕ is associated with $[\![\mho^\omega]\!]^\top(!(\top)(\delta\phi))$.[63] Without getting into the details (see Segerberg (1994) for those), the semantics yields the following requirement conditions for this formula (the permission aspect being redundant, we ignore it for the sake of simplicity):

(165) $\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} [\![\mho^\omega]\!]^\top(!(\top)(\delta\phi))$ iff
$\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} [\![\mho]\!]^{\top k}(!(\top)(\delta\phi))$, for all $k \in \mathbb{N}$ iff

$\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} !(\top)(\delta\phi)$ &
$\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} [\![\mho]\!]^\top(!(\top)(\delta\phi))$ &
$\mathcal{M},c,w \Vdash_{\mathcal{L}_{ILA}} [\![\mho]\!]^\top([\![\mho]\!]^\top(!(\top)(\delta\phi)))$ &... iff

$\bigcap t_c(a_c)(w) \cap i_c \subseteq [\![\delta\phi]\!]_{\mathcal{M},c}$ &
$\forall\langle v,v'\rangle \in \bigcap t_c(a_c)(w) \cap i_c : \bigcap t_c(a_c)(v') \cap i_c \subseteq [\![\delta\phi]\!]_{\mathcal{M},c}$ &
$\forall\langle v,v'\rangle \in \bigcap t_c(a_c)(w) \cap i_c : \forall\langle u,u'\rangle \in t_c(a_c)(v') \cap i_c : \bigcap t_c(a_c)(u') \cap i_c \subseteq [\![\delta\phi]\!]_{\mathcal{M},c}$ &...

These are precisely the requirement conditions envisioned above for stable commands.

The envisaged contrast between ephemeral and stable commands (as with the contrast between order-sensitive and order-insensitive commands) has interesting consequences for axiomatizing the deontic side of the logic. Consider the following $\mathcal{L}_{DLA}$ analogues for the **D4** axiom of standard deontic logic, $O\phi \rightarrow OO\phi$, which enforces transitivity of the

62. For a proper treatment of an infinitary propositional dynamic logic with the Kleene Star, see, e.g., Segerberg (1994).

63. Thus *always* is interpreted in the simplest case as a scope-taking dynamic modal operator $[\![\mho^\omega]\!]^\top(\cdot)$. In cases of conditional stable commands—e.g., (55a)—the restriction argument of the operator may be something other than $\top$.

standard deontic accessibility relation.[64] Note that these axioms are equivalent, by result (134).

(166) $O(\pi)(\delta\phi) \rightarrow O(\pi)(\delta O(\pi)(\delta\phi))$

(167) $[\![\mho]\!]^{\pi}(\phi) \rightarrow [\![\mho]\!]^{\pi}([\![\mho]\!]^{\pi}(\phi))$

But there is a problem with these axioms. The former (and, so, the latter) implies that obligations (and thereby commands) *never lapse*: commands that are in force at w remain in force at the terminal state of each w-preferred execution of them.

Proof. Suppose $\mathcal{M}, c, w \Vdash_{\mathcal{L}_{DLA}} O(\pi)(\delta\phi)$.
Then $\mathcal{M}, c, w \Vdash_{\mathcal{L}_{DLA}} O(\pi)(\delta O(\pi)(\delta\phi))$.
So $\forall \langle v, v' \rangle \in \mathit{min}[i_c \cap ([\![\pi]\!]_{\mathcal{M},c} \times W), \leq_{t_{c(a_c)(w)}}] : \langle v, v' \rangle \in [\![\delta O(\pi)(\delta\phi)]\!]_{\mathcal{M},c}$.
But then $\mathcal{M}, c, v' \Vdash_{\mathcal{L}_{DLA}} O(\pi)(\delta\phi)$.

But it is the essence of ephemeral commands to lapse upon fulfillment. It follows that these axioms must be treated as anathema in a logic that represents the distinction between ephemeral and stable commands as we choose to here. As was the case with the $\mathcal{L}_{DLA}$ analogue of **OU**, this is somewhat surprising. The failure of relatively standard axioms of SDL to carry over to our system is principally a consequence of the temporality implicitly built into the logic.

4.13 Rapprochement and Transition to Update Semantics

The semantics for imperatives developed in this section of the paper has it that imperatives express both constraints on and permissions for certain kinds of planning behavior of a subject of authority (what we have been referring to as an "addressee") at a context. Specifically: an imperative is required (or in force) at a context just in case the To-Do List of the subject of authority—the set of constraints that impinge on her planning behavior—satisfies a certain complex relation—the most sophisticated statement of which was to be found in (137)—with respect to the information that is relevant at the context. As such, it is most natural to regard it as an implementation of the motivations that underlie a logic of planning. While natural, however, the choice is unforced. Resisting the choice will reveal a way in which a single formalism can serve as an implementation of the intuitions that underlie both logics of content and logics of planning. It also provides a nice way of connecting the ostensibly static PDLA framework we have been laboring to develop to an update semantics for imperatives: these too are also, to an extent, two sides of what is roughly the same coin.

The logic of content developed in the prior section of the paper proceeded from the intuitions that (i) imperative logic should be normative for endorsement by an authority, and (ii) what constrains endorsement of an imperative the agent's desires, together with the information at a context in which the agent is an authority—the issuer or "speaker". This led us to develop a logic of imperatives in terms of a deontic modal logic designed to track when imperatives were suitable expressions of the authority's desires, in light of the relevant information, and when they were not. But there are other options. Indeed, a To-Do List semantics can capture both of the intuitions underlying a logic of content.

The To-Do List semantics understands imperative logic in terms of practical or planning content: imperatives express constraints on an agent's planning in a situation (as well

64. Transitivity is usually viewed as a desirable property of a deontic accessibility relation. See Chellas (1980) for a defense (with reservations) and Vorobej (1982) for an lucid argument against those reservations.

as permissions). As such, it seems almost essentially subject- (rather than authority-) or addressee- (rather than speaker-) oriented. But, of course, constraints on her subject's planning (as well as the occasional freedom) are *precisely* what an authority seeks to impose with an imperative. Put differently, the *object* of the authority's endorsement is naturally understood as being exactly what the logic of planning regards as the practical content of the imperative: the constraints on and permissions for a subject's planning that it expresses. Rather than identifying the content of an imperative with an expression of the desires of the authority, we might just as well identify it with its practical content. An authority intuitively should endorse an imperative just when its practical content is a suitable expression of her desires (although there is, at this juncture, nothing demanding explicit representation of her desires in the logic). One imperative follows from another just when endorsing the practical content of the latter commits an authority to endorsing the practical content of the former. It would seem, then, that logics of content and logics of planning express different, yet formally compatible, perspectives on the very same phenomenon. Logics of content take as their subject matter the will of the authority *as it concerns the plans of the subject*, while logics of planning take as their subject matter the plans of the subject *in light of the will of the authority*.

Since there is no reason to think that a single formalism cannot do justice to both perspectives, it make sense to compare logics that emanate from these different perspectives (whereas formerly we would have regarded them as incommensurable). As we've seen, the logic of planning, such as it is, that we developed in this section of paper handles a range of interesting phenomena about the imperative that the logic of content, such as it was, developed in the paper's prior section could not. This constitutes a strong reason for preferring the former, and for discarding the latter.

This is all pretty vague, so let us try to fill in the sketch a bit. (A bonus: doing so will give us a nice lead-in to our discussion of update semantics.) Let us understand the *static content* of a non-modal imperative formula π in a context c (with respect to model $\mathcal{M}$) as the characteristic function $\chi_c(\pi) : \mathcal{T} \mapsto \{0,1\}$ of a set $\mathsf{inad}_{\chi_c}(\pi)$ of inadmissible To-Do Lists—those mapped to 1 by $\chi_c(\pi)$—for an addressee.[65] What I have in mind is, ignoring permission for sake of simplicity, something like the following:

(168) Where $w \in \{w \mid \exists w' : \langle w, w' \rangle \in i_c\}$:
$\mathsf{inad}_{\chi_c}(!(\phi)(\alpha)) = \{\tau_{c,a_c} : W \mapsto 2^{\mathcal{A}} \mid \mathit{min}[i_c \cap (\llbracket \phi \rrbracket_{\mathcal{M},c} \times W), \leq_{\tau_{c,a_c}(w)}] \nsubseteq \llbracket \alpha \rrbracket_{\mathcal{M},c}\}$

The formula $!(\phi)(\alpha)$ expresses a function $\chi_c(!(\phi)(\alpha))$ according to which a To-Do List for the addressee τ_{c,a_c} is admissible just in case at any input state w compatible with the relevant information, $\tau_{c,a_c}(w)$ yields an obligation to perform α in the best ϕ-initial situations. Informally, $!(\phi)(\alpha)$ expresses a constraint on a subject's To-Do List: a To-Do List counts as admissible just in case, from its vantage, α is required if ϕ. Suppose we understand requirement conditions for imperatives as relative to To-Do Lists directly, rather than indirectly (via contexts). Then we have $\mathsf{inad}_{\chi_c}(!(\phi)(\alpha))$ being just the set of To-Do List functions τ_{c,a_c} such that $\mathcal{M}, \tau_{c,a_c}, w \Vdash_{\mathcal{L}_{ILA}} !(\phi)(\alpha)$, for any input state w compatible with the relevant information.

The notion that imperative formulas of $\mathcal{L}_{ILA}$ might be treated as expressing such characteristic functions is a useful one for reconciling logics of content and logics of planning. Issuers of imperatives (authorities) might be represented as expressing some sort of pro-

65. Eric Swanson calls this sort of approach—on which the semantic value of a formula is typed as the characteristic function of a set of inadmissible types of cognitive state—"constraint semantics." See Swanson (2008a) for an short development of the idea, and Swanson (2006: Ch. 2) for a detailed outline of a constraint semantics for a fragment of natural language.

attitude toward such a function: an attitude of endorsement or approval of that function—or, more precisely, that function's representation of the normative state of play. Endorsing such a function commits an authority to endorsing another just when the pre-defined imperative entailment relation holds between the relevant formulas. Such functions might also be thought to constrain the planning apparatus of their addressees: an agent whose planning apparatus fails to be representable with an admissible To-Do List—a To-Do List not in $\mathsf{inad}_{\chi_c}(\pi)$—violates the constraint expressed by the imperative formula. To be cooperative, she must assimilate it, by modifying her To-Do List so that it is no longer describable as inadmissible, by the lights of $\mathsf{inad}_{\chi_c}(\pi)$.

The connection of static content to update semantics is apparent: update semantics could be understood as stemming from the impulse to model this sort of assimilation. Supposing that an authority uses an utterance to express endorsement of a characteristic function χ and that her subject's antecedent planning apparatus is not representable with an admissible, by the lights of χ, To-Do List, being cooperative will entail modifying her To-Do List, so that it becomes so representable. Assigning *update potentials*—functions from cognitive states into updated cognitive states of the same type—to formulas of our language is thus treated as a matter of describing a function that maps inadmissible cognitive states into suitable admissible ones (and describing what the updated states look like), and admissible cognitive states into themselves. We now turn our attention to this task.

5 DYNAMICS AND DYNAMIC SEMANTICS

The central problem in giving a pragmatic/dynamic analysis of imperatives is accounting for their peculiar force—their essentially *performative* character. In principle, it is possible to examine this question without either (i) committing oneself to a definite view about the update (or context change) potential of imperative formulas or (ii) trying to leverage a view of update potentials to define dynamic analogues of static notions of satisfaction, requirement, entailment, etc. Addressing the questions in which we are interested does, however, require that we do both.

Our study of imperative dynamics begins in a rather different place from most recent work on the subject. We have a complex apparatus already in place, and, rather than tearing it down and building it up piecemeal, we will simply presuppose a good deal of it as we build a dynamics for imperatives. This has its advantages—a complicated apparatus turns out to be capable of handling many facts about the pragmatics of imperatives that simpler accounts struggle with. Nevertheless, as will become clear, in spite of large differences in apparatus, the basic intuition behind our account is a familiar one: the performative force of imperatives will be accounted for by construing them as in the business of updating the To-Do Lists of their addressees.

This section is structured around building a dynamics for imperatives that secures certain desiderata that an analysis of imperatives' performative force ought to have. Subsections are generally devoted to accounting for one desideratum in particular. Once a reasonably adequate apparatus for handling the relevant dynamic phenomena is in place, we move on to define a dynamic analogue of requirement conditions (and imperative entailment) for imperative formulas. We close with a comparison of our dynamic system to the static system developed in the prior section and a polemic of sorts against recent attempts to marginalize non-dynamic semantic approaches to the imperative.

5.1 Performative Force

Before beginning in earnest, it is important to state in a reasonably precise way (1) what it is we mean when we say that an imperative (or, more precisely, an imperative utterance) has performative force and (2) the properties a formal analysis should exhibit in order to count as a satisfactory account of this phenomenon. Regarding (1), I take it that by claiming that imperatives have performative force, we mean that utterances of imperatives in natural language (of the sort that we use imperative formulas of the imperative object language $\mathcal{L}_{ILA}$ to represent) (i) function to *introduce some species of obligation or requirement on their addressees*, and (ii) fulfill this function *reliably*. Regarding (2), it would therefore seem that providing a satisfactory analysis of imperative force will require modeling imperative *discourse* in our formal system—in particular, devising a dynamic formal system that is suited to representing the impact of utterances in natural language on a discourse. I will suppose that the best way to do this is by defining new interpretations for formulas in the imperative object language—update potentials—which take a context as input and return an updated context as output. An assignment of update potentials to formulas of the language succeeds insofar as it represents imperative utterances as reliably imposing the relevant requirements and obligations on their addressees.

What does it mean to say that imperative utterances reliably impose requirements and obligations on their addressees? I propose we cash this out by representing imperative utterances as affecting the truth-values of deontic formulas—in particular, causing certain descriptions of obligation go from false (at the prior context) to true (at the updated context)—and requirement-values of imperative formulas—in particular, making it the case that certain imperatives go from failing to be in force (at the prior context) to being in force (at the updated context). That is to say: updating a context according to an imperative reliably yields a new context in which certain deontic statements are true and certain imperatives are in force. The formal import for the dynamic semantics is this: we will seek to define an update potential for the imperative that reliably *alters contextual parameters* relevant to evaluating whether relevant deontic formulas are satisfied and imperative formulas are in force.

Concretely: suppose that I am ordered to write a paper about imperatives at a context c. Then updating c with the order should yield a new context c^* where, inter alia, I must (in view of the constraints impinging on my planning) do so and the imperative *Write a paper about imperatives!*—or, rather, the $\mathcal{L}_{ILA}$ representation we assign to it—is in force. It is, at the same time, clearly possible that, at the prior context c, I needn't do so, since, at c, *I haven't been told to*. A bit more formally: suppose $!(\phi)(\alpha)$ is uttered at c, and let c^* be c updated with $!(\phi)(\alpha)$. When update succeeds, a good account of imperative dynamics should yield the following predictions.

(169) $c^* \Vdash O(\phi)(\alpha)$, although, possibly, $c \nVdash O(\phi)(\alpha)$

(170) $c^* \Vdash !(\phi)(\alpha)$, although, possibly, $c \nVdash !(\phi)(\alpha)$

Three notes on notation. (i) Since we have finished toying around with the object languages, we will avoid clutter by avoiding labeling the symbols for satisfaction, requirement, and entailment relations in the metalanguage, except where necessary. (ii) Since our focus is now on contexts, we can (and will) generally leave the role of models implicit, without any loss of information. (iii) We use a new semantic relation, $\Vdash$, to denote validity in a context—i.e., satisfaction or requirement with respect to all of the possibilities relevant at the context: $c \Vdash \phi$ shall abbreviate either $\forall w \in \{w \mid \exists w' : \langle w, w' \rangle \in i_c\} : \mathcal{M}, c, w \vDash \phi$ or $\forall w \in \{w \mid \exists w' : \langle w, w' \rangle \in i_c\} : \mathcal{M}, c, w \Vvdash \phi$, as context will make clear.

It is important to note that (169) and (170) can come apart, on account of the permissive dimension of imperatives: the practical content of an imperative is not, we have argued, exhausted by an expression that so-and-so action is *required* in so-and-so situation, but will often—e.g., in the case of choice-offering disjunctive imperatives—involve an expression of permission that cannot be "derived" from the command content of the imperative. For now, however, we will focus on the introduction of requirements and constraints—predicting (169)—rather than permissions, saving a discussion of the permissive dimension of imperative force for a bit later on. We will build our dynamic analysis in two steps. The first will focus on obligation introduction, the second on permission introduction. Putting the pieces together yields an account that predicts (170).

As already noted, validating any of the relevant facts demands defining dynamic notions—in particular, a update potential for imperative formulas of the imperative object language—that predict them. Insofar as obligations are concerned, there are two (and only two) parameters of the context that are relevant to deciding whether or not an imperative or deontic formula is validated there—the contextual modal base and the contextual ordering-source (To-Do List). To yield the desired prediction, an update potential for the imperative must operate on one (possibly both) of these parameters.

5.2 Updating i_c

We have given a bit of the game away in making it clear that we would be endorsing an update potential for the imperative that operates on the ordering-source (or parameters related to the ordering-source). But there is, to be sure, nothing crazy about thinking that the performative force of imperatives is to be accounted for in terms of the addition of information to the modal base. Indeed, for a modal analysis of the imperative, alteration of the modal base is a natural place to start: imperatives might, for example, add the semantic value of an obligation-describing modal to the modal base.[66]

A toy implementation of this basic idea for the imperative object language $\mathcal{L}_{ILA}$. We define update potentials for the non-modal fragment of $\mathcal{L}_{ILA}$.[67] We will have $!(\phi)(\alpha)$ altering the modal base by adding the information in $[\![O(\phi)(\alpha)]\!]_c$ to update i_c's representation of the current state of the world. Non-imperative formulas will simply add the information they express about the current state of the world to i_c. Let $\|\cdot\|_{\mathcal{M}}$ be a function from formulas of $\mathcal{L}_{ILA}$ into functions from contexts of $\mathcal{M}$ into contexts of $\mathcal{M}$ (we will hereafter omit mention of models), and model a context c as before: as the ordered tuple $\langle s_c, a_c, i_c, t_c\rangle$ (ignoring non-derivative rights lists for now). The relevant updates will be given as follows:

(171) Where ϕ does not contain an imperative:
$c\|\phi\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*}\rangle$
- $s_{c^*} = s_c$, $a_{c^*} = a_c$, $t_{c^*} = t_c$
- $i_{c^*} = i_c \cap ([\![\phi]\!]_c \times W)$

(172) $c\|!(\phi)(\alpha)\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*}\rangle$
- $s_{c^*} = s_c$, $a_{c^*} = a_c$, $t_{c^*} = t_c$
- $i_{c^*} = i_c \cap ([\![O(\phi)(\alpha)]\!]_c \times W)$

This approach, should it succeed, would have two clear virtues. First, of course, it predicts automatically that $c^* \Vdash O(\phi)(\alpha)$. And, second, it makes clear how handling the static semantics of imperatives in terms of the semantics of deontic modals might be connected to a theory of imperative force. Performative force is achieved (putatively) by adding the semantic value of the deontic formula whose satisfaction conditions give the imperative's requirement conditions to the modal base's depiction of the present state of the world.

66. Schwager, for her part, seems to understand performative force in terms of update of the Common Ground: imperatives update contexts by adding the information *about the ordering-source* (in her case, the speaker's desires) to the Common Ground (see Schwager 2006: 10). It is not clear from her presentation what information about the desires of the speaker has to do with performative force.

67. One immediate problem with this approach is that there is no clearly sensible way of defining update potentials for formulas of the form $[\beta](!(\phi)(\alpha))$ or $[\![\beta]\!]^{\pi}(!(\phi)(\alpha))$. But I won't harp on this here.

While we were dubious that the non-integrability of a static semantics for imperatives with a theory of imperative force constituted any real objection to such a semantics (§3.3), we can appreciate the point that such integrability is a praiseworthy (if not obligatory) feature of a semantics.

There are, however, empirical and conceptual difficulties. Empirically, we do not have an account of performative force, because while we have an account that reliably predicts that obligations are in force at the updated context, this comes at the expense of a genuine account of obligation *introduction*. Consider a schematic case. We begin by making explicit some assumptions that we have to this point left implicit: (i) contexts are properly understood as "inhabitants" of worlds, so that a context c supplies information about the world w_c locating c;[68] (ii) updating the modal base parameter of the context does *not* alter the state in which c is located; (iii) modal bases are realistic, in the sense that $w_c \in \{w \mid \exists w' : \langle w, w' \rangle \in i_c\}$. The assumption that $w_c \nvDash O(\phi)(\alpha)$ leads immediately to contradiction.

> Proof. Let $w_c \nvDash O(\phi)(\alpha)$. Since $w_{c^*} = w_c$, $w_{c^*} \nvDash O(\phi)(\alpha)$. But, by supposition, $w_{c^*} \in \{w \mid \exists w' : \langle w, w' \rangle \in i_{c^*}\}$. So $c^* \nVdash O(\phi)(\alpha)$. Contradiction.

But it should, of course, be *possible* that $w_c \nvDash O(\phi)(\alpha)$.

This leads into the conceptual difficulties with the view, which are twofold. First, it is strange for $!(\top)(\alpha)$ to update i_c by, in effect, inducing the presupposition that α is required. It is entirely ordinary to suppose that at i_c it is presupposed that the subject is under no obligation to perform α—i.e., $c \Vdash \neg O(\phi)(\alpha)$, which implies that $w_c \nvDash \neg O(\phi)(\alpha)$. Imperatives are in the business of introducing *new* obligations (and perhaps also canceling prior presuppositions about obligations), including (indeed, perhaps especially) obligations antecedently presupposed not to be in place. But in any case where $c \Vdash \neg O(\phi)(\alpha)$, updating i_c with $!(\top)(\alpha)$ will land us in broken context, with an empty modal base. Second, the pragmatic force of an *assertion* that π is usually modeled as an addition of the information that π to a modal base's representation of the state of the world (Stalnaker 1978, 2002). But this is precisely how the proposal under consideration construes the performative force of an imperative. The distinction between assertive force—the sort of force traditionally associated with modal base update—and performative force is blurred.

The empirical and conceptual difficulties for this approach to imperative dynamics make a strong case for a different tack. We begin to develop one in the following section.

5.3 Updating t_c

Another option—the one we will pursue—is to construe imperatives as To-Do List (ordering-source) updaters.[69] Rather than defining update potentials exclusively for the non-modal fragment of the imperative object language, we define update potentials for all of $\mathcal{L}_{ILA}$.

68. This is a natural and useful assumption. Others that make it include Isaacs & Potts (2003); Potts (2003). We have left this implicit in the interest of formal simplicity: the world parameter of a context is semantically idle.

69. Isaacs & Potts (2003); Mastop (2005); Portner (2004, 2008); Potts (2003); Veltman (2008) are the major references for this type of view in contemporary linguistics. The overall shape of the view given here is most indebted to the formulation in Portner (2008). Giving originality its due, Segerberg (1990)'s notion of an action-guiding Command System is *quite* similar to the contemporary notion of a To-Do List (although Segerberg does not himself attempt to give an account of imperative dynamics). Indeed, the account he builds around the notion of a Command System is rather more sophisticated than contemporary accounts, in that it (i) construes the relevant parameter as sets of actions, rather than propositions, properties, or the like; (ii) is aware that the parameter has a central role to play in giving an account of the Ross Paradox; and (iii) is adaptable to handling temporal phenomena about the imperative, of the sort we examined above.

That is to say, $\|\cdot\|_{\mathcal{M}} : \mathcal{L}_{ILA} \mapsto \{f \mid f : \mathcal{C} \mapsto \mathcal{C}\}$ is a total function defined for every formula of $\mathcal{L}_{ILA}$.

The easiest way of doing this properly requires a small modification in our understanding of To-Do Lists, somewhat along the lines of the neighborhood semantics given in §3.5.1. We allow proposition-action pairs to occur on a To-Do List with respect to an index of evaluation, with the following interpretation:

(173) $\langle p, \Phi\rangle \in t_c(d)(w) \Leftrightarrow \forall\langle v, v'\rangle \in min[i_c \cap (p \times W), \leq t_c(d)(w)] : \Phi \in t_c(d)(v)$

The "presence" of $\langle p, \Phi\rangle$ on a To-Do list at w means that the To-Do List at the initial state of each best-from-w p-initial transition contains Φ.[70] Note the following consequence of this definition: we may, in effect, replace the ordering-source semantics for deontic O with a straight neighborhood semantics.

(174) $c, w \vDash O(\phi)(\alpha)$ iff $\langle[\![\phi]\!]_c, [\![\alpha]\!]_c\rangle \in t_c(a_c)(w)$

Proof. $c, w \vDash O(\phi)(\alpha)$ iff (by the neighborhood semantics of §4.10.3)
$\forall\langle v, v'\rangle \in min[i_c \cap ([\![\phi]\!]_c \times W), \leq_{t_c(a_c)(w)}] : [\![\alpha]\!]_c \in t_c(d)(v)$ iff
$\langle[\![\phi]\!]_c, [\![\alpha]\!]_c\rangle \in t_c(a_c)(w)$

(175) $c \Vdash O(\phi)(\alpha)$ iff $\forall w \in \{w \mid \exists w' : \langle w, w'\rangle \in i_c\} : \langle[\![\phi]\!]_c, [\![\alpha]\!]_c\rangle \in t_c(a_c)(w)$

Proof. Immediate from definition of $\Vdash$.

We proceed to define update potentials for formulas of $\mathcal{L}_{ILA}$ as follows. Informal explanations are appended after each definition, as well as a short proof that the proposed update potentials secure the desired results about performative force.

5.3.1 Non-Imperative Formulas

Update with ordinary formulas is still treated as restriction of the modal base.

(176) Where ϕ does not contain an imperative, $c\|\phi\|$ is as defined in (171).

5.3.2 Imperative Formulas

It is fairly easy to define update with an imperative formula $!(\phi)(\alpha)$ so that the updated context validates $O(\phi)(\alpha)$. We do so as follows.

(177) $c\|!(\phi)(\alpha)\|$ is defined iff $t_{c^*}(a_c)$ (defined below) satisfies **CON***, **AV2****, **KE2***.
When defined, $c\|!(\phi)(\alpha)\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*}\rangle$
- $s_{c^*} = s_c$, $a_{c^*} = a_c$, $i_{c^*} = i_c$
- $t_{c^*}(d) = t_c(d)$, for all $d \neq a_c$
- $t_{c^*}(a_c) = t_c(a_c)\ [w\ /\ t_c(a_c)(w) \cup \{\langle[\![\phi]\!]_c, [\![\alpha]\!]_c\rangle\}]$,

for all $w \in \{w \mid \exists w' : \langle w, w'\rangle \in i_c\}$

Update is undefined when it would yield an broken To-Do List.[71] When update is defined, a_c's updated To-Do List function is such that for any state w over which it is defined, w

70. If the reader finds this confusing, imagine To-Do Lists on which proposition-action pairs occur as distinct from "actual" To-Do Lists—those containing only actions—and treat the former as a construction built out of the latter, or vice versa.

71. **CON***, **AV2****, and **KE2*** are understood to apply to "atomic" actions, not proposition-action pairs.

contains $\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle$. I.e., the initial state v of any best-from-w, ϕ-initial transition is such that a_c's To-Do List at v contains $[\![\alpha]\!]_c$. We have the following result.

(178) $c\|!(\phi)(\alpha)\| \Vdash O(\phi)(\alpha)$, when $c\|!(\phi)(\alpha)\|$ is defined

Proof. Immediate given (174), (175), $i_{c^*} = i_c$, $[\![\phi/\alpha]\!]_{c^*} = [\![\phi/\alpha]\!]_c$.

5.3.3 Basic Dynamic Formulas

The formula $[\beta](!(\phi)(\alpha))$ expresses a restricted command: conditional on the performance of β, the subject is required to perform α if ϕ. We thus define update with the dynamic modal formula $[\beta](!(\phi)(\alpha))$ so that the updated context validates $[\beta](O(\phi)(\alpha))$.

(179) $c\|[\beta](!(\phi)(\alpha))\|$ is defined iff $t_{c^*}(a_c)$ satisfies **CON***, **AV2****, **KE2***.
When defined, $c\|[\beta](!(\phi)(\alpha))\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*} \rangle$

- $s_{c^*} = s_c$, $a_{c^*} = a_c$, $i_{c^*} = i_c$
- $t_{c^*}(d) = t_c(d)$, for all $d \neq a_c$
- $t_{c^*}(a_c) = t_c(a_c)\ [v\ /\ t_c(a_c)(v) \cup \{\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle\}]$, $v \in \{v \mid \exists w : \langle w, v \rangle \in i_c \cap [\![\beta]\!]_c\}$

a_c's updated To-Do List function is such that for any state v that is the terminal state of an i_c-compatible transition of β, a_c's To-Do List at v contains $\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle$. For all other states not satisfying this condition, the value of the function is unchanged. The function of the modal operator $[\beta]$ is, in effect, to shift the context (by shrinking the modal base to the set of β-transitions) and update the result with $!(\phi)(\alpha)$. We thus have the following result.

(180) $c\|[\beta](!(\phi)(\alpha))\| \Vdash [\beta](O(\phi)(\alpha))$, when $c\|[\beta](!(\phi)(\alpha))\|$ is defined

Proof. Again immediate given (174), (175), $i_{c^*} = i_c$, $[\![\phi/\alpha/\beta]\!]_{c^*} = [\![\phi/\alpha/\beta]\!]_c$, and the semantics of $[\beta]$.

5.3.4 Complex Dynamic Formulas

The formula $\lceil\!\lceil \beta \rceil\!\rceil^{\pi}(!(\phi)(\alpha))$ expresses a doubly restricted command: the subject is required to perform α if ϕ, supposing a good-enough, ϕ-initial execution of β. We thus define update with $\lceil\!\lceil \beta \rceil\!\rceil^{\pi}(!(\phi)(\alpha))$ so that the updated context validates $\lceil\!\lceil \beta \rceil\!\rceil^{\pi}(O(\phi)(\alpha))$.

(181) $c\|\lceil\!\lceil \beta \rceil\!\rceil^{\pi}(!(\phi)(\alpha))\|$ is defined iff $t_{c^*}(a_c)$ satisfies **CON***, **AV2****, **KE2***.
When defined, $c\|\lceil\!\lceil \beta \rceil\!\rceil^{\pi}(!(\phi)(\alpha))\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*} \rangle$

- $s_{c^*} = s_c$, $a_{c^*} = a_c$, $i_{c^*} = i_c$
- $t_{c^*}(d) = t_c(d)$, for all $d \neq a_c$
- $t_{c^*}(a_c) = \mathsf{cl}^n[t_c(a_c)]$, the result of executing $\mathsf{cl}[\cdot]$ on $t_c(a_c)$ n times, where n is the smallest k such that $\mathsf{cl}^k[t_c(a_c)] = \mathsf{cl}^{k-1}[t_c(a_c)]$.

  $\mathsf{cl}[t_c(a_c)] =_{df} t_c(a_c)\ [v'\ /\ t_c(a_c)(v') \cup \{\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle\}]$,
  for all $v' \in \{v' \mid \exists w \exists v : \langle v, v' \rangle \in \mathit{min}[i_c \cap ([\![\pi]\!]_c \times W) \cap [\![\beta]\!]_c, \leq_{t_{c^*}(a_c)(w)}]\}$

This definition has it that a_c's updated To-Do List function is such that for any state v' that is the terminal state of an i_c-compatible, π-initial, best-at-w transition of β (for some w

compatible with i_c), a_c's To-Do List at v' contains $\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle$. We will prove this.

(182) Claim. Let $w \in \{w \mid \exists w' : \langle w, w' \rangle \in i_{c^*}\}$. Then:
$\langle v, v' \rangle \in \textit{min}[i_{c*} \cap ([\![\pi]\!]_{c^*} \times W) \cap [\![\beta]\!]_{c^*}), \leq_{t_{c^*}(a_{c^*})(w)}] \Rightarrow \langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle \in t_{c^*}(a_{c^*})(v')$

Proof. Let $\langle v, v' \rangle \in \textit{min}[i_{c*} \cap ([\![\pi]\!]_{c^*} \times W) \cap [\![\beta]\!]_{c^*}), \leq_{t_{c^*}(a_{c^*})(w)}]$.

We know that for all u' such that for some $w, u \in \{w \mid \exists w' : \langle w, w' \rangle \in i_{c^*}\}$, $\langle u, u' \rangle \in \textit{min}[i_{c^*} \cap ([\![\pi]\!]_{c^*} \times W) \cap [\![\beta]\!]_{c^*}, \leq_{t_{c^*}(a_{c^*})(w)}]$ implies $\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle \in t_{c^*}(a_{c^*})(u')$, unless $\mathsf{cl}[t_{c^*}(a_{c^*})] \neq t_{c^*}(a_{c^*})$, which is impossible.

Since $w, v \in \{w \mid \exists w' : \langle w, w' \rangle \in i_{c^*}\}$ and $\langle v, v' \rangle \in \textit{min}[i_{c*} \cap ([\![\pi]\!]_{c^*} \times W) \cap [\![\beta]\!]_{c^*}), \leq_{t_{c^*}(a_{c^*})(w)}]$, the claim follows.

It follows that:

(183) $c \| [\![\beta]\!]^\pi(!(\phi)(\alpha)) \| \Vdash [\![\beta]\!]^\pi(O(\phi)(\alpha))$, when $c \| [\![\beta]\!]^\pi(!(\phi)(\alpha)) \|$ is defined

Proof. Again immediate given (174), (175), $i_{c^*} = i_c$, $[\![\phi/\pi/\alpha/\beta]\!]_{c^*} = [\![\phi/\pi/\alpha/\beta]\!]_c$ and the semantics of $[\![\beta]\!]^\pi$.

5.3.5 Conceptual Virtues

To sum up: our definitions of update potentials for formulas of $\mathcal{L}_{ILA}$ predict that, when updates are defined, updating a context with an imperative formula of any stripe reliably introduces the proper sorts of obligations on their subjects. But this is not all. A nice additional bonus: construing imperative formulas or formulas containing imperative formulas as To-Do List updaters predicts that and explains how imperative and "assertive" force—the sort of force associated with updating a context with a non-imperative formula of the language—are distinct. A formula (or, more precisely, an utterance thereof) has assertive force (may be used assertively) just in case its update potential operates on the informational parameter of the context, imperative force just in case its update potential operates on the action-guiding parameter. Understanding imperative dynamics in terms of update of the modal base founders on this distinction.

As we noted above, one putative virtue of the imperatives-as-modal-base-updaters approach is that there is a real (indeed, obvious) connection between the static semantics and the theory of imperative force. The account on offer shares this property. Let us generalize the idea of constraint semantics (§4.13) to formulas of $\mathcal{L}_{ILA}$ besides those instantiating the form $!(\phi)(\alpha)$. We associate a formula π of $\mathcal{L}_{ILA}$ with a function $\chi_c(\pi) : \mathcal{T} \mapsto \{0,1\}$ of a set characterizing the set $\mathsf{inad}_{\chi_c}(\pi)$ of inadmissible To-Do Lists—To-Do Lists that *fail to validate* π. (Note on notation: $\tau_{c,a_c} \Vdash \pi$ iff π is required at every $w \in \{w \mid \exists w' : \langle w, w' \rangle \in i_c\}$, if we utilize τ_{c,a_c} as the ordering source.)

(184) $\mathsf{inad}_{\chi_c}(\pi) = \{\tau_{c,a_c} : W \mapsto 2^{\mathcal{A}} \mid \tau_{c,a_c} \nVdash \pi\}$

Let π be any formula of $\mathcal{L}_{ILA}$ that is or contains an imperative. In every case, the update potential defined above for π does nothing more (and nothing less) than to ensure that a post-update To-Do List validates π. That is to say: when $t_c(a_c) \notin \mathsf{inad}_{\chi_c}(\pi)$, $c\|\pi\| = c$; and when $t_c(a_c) \in \mathsf{inad}_{\chi_c}(\pi)$, $c\|\pi\| = c^*$, where $t_{c^*}(a_{c^*}) \notin \mathsf{inad}_{\chi_c}(\pi)$. Genuine, non-vacuous update is triggered when, and only when, τ_{c,a_c} is inadmissible at c—i.e., $\tau_{c,a_c} \nVdash \pi$.[72] In

72. All these facts are easy to verify from things proved in the prior section. I will not bother with the proofs.

a nutshell: the static side of the semantics supplies synchronic constraints on To-Do Lists, while the dynamic side models diachronic satisfaction of these constraints.

5.4 *Permissive Force*

We understand the permissive force of imperatives[73] as being roughly conceptually analogous to their obligation-imposing force: utterances of imperatives in natural language (of the sort represented with formulas of $\mathcal{L}_{ILA}$) function to confer rights, freedoms, or entitlements on their addressees, and they fulfill this function reliably. The implementation of this idea, however, will differ significantly from our implementation of the corresponding intuition about obligation-imposing force. To illustrate: while the fact that $c \Vdash \neg O(\phi)(\alpha)$ is no obstacle, per se, to having it be the case that $c\|!(\phi)(\alpha)\| \Vdash O(\phi)(\alpha)$, the fact that $c \Vdash \neg P(\phi)(\alpha)$ certainly ought to be an obstacle to having it be the case that $c\|!(\phi)(\alpha)\| \Vdash P(\phi)(\alpha)$. While imperatives can cancel prior permissions—within certain bounds (which we will explore below)—we will suppose that the permissions they express *must not conflict with prior obligations*.[74] So, the technical sense in which imperatives introduce rights is distinct from the sense in which they introduce obligations: they do not *alter* the truth-values of deontic formulas that describe the permissions they express. Indeed, if a context c is updatable with an imperative formula, the permissions that the formula expresses must already be in force at c. There is, nevertheless, an obvious sense in which imperatives confer or introduce rights that their subjects did not have before. It is just that, given our assumptions, we cannot cash this intuition out in terms of alteration of contextual parameters with an aim to establishing the truth of relevant permission formulas. Instead, I suggest we cash it out in terms of addition to a Non-Derivative Rights List (cf. §4.11.2).

Very roughly: we will conceive of a NDRL function for an individual d, $r_c(d)$, as a function from worlds to sets of *actions that the agent has been granted permission to perform in those worlds* (rights, entitlements, freedoms). Actions that an agent has been granted permission to perform are distinguished from merely permissible actions: an action is permissible just in case the relevant deontic formula of $\mathcal{L}_{DLA}$ expressing that the action is permissible is true. An action is a right just in case the relevant deontic formula of $\mathcal{L}_{DLA}$ expressing that the action is permissible is true, *and*, additionally, it occurs on the agent's NDRL. NDRLs, then, characterize sets of privileged permissions at a world. Imperatives, which are in the business of granting permissions, are conceived as adding the permissions they express to an agent's rights. The function of an NDRL is, as intimated above, essentially dynamic: the update semantics uses them to check whether updating the context with an imperative formula would conflict with any of the agent's rights. While imperatives can cancel permissions, we will assume that they cannot cancel rights. Update with an imperative is rejected whenever either (i) the imperative attempts to cancel any of the subject's rights or (ii) the updated NDRL function for the subject would fail to satisfy any of the constraints introduced on NDRLs in §4.11.2, repeated here for convenient reference.

(185) If $\Phi \in r_c(d)(w)$, then $\Phi \in r_c^-(d)(w) = \{\Phi \mid \overline{\Phi} \notin t_c(d)(w)\}$.

(186) If $\Phi \in t_c(d)(w)$, then $\Phi \in r_c(d)(w)$ **[CR]**

(187) If $\Phi_1 \in r_c(d)(w)$, then $\Phi_1 \cap \Phi_2 = \emptyset$ implies $min[i_c, \leq_{t_c(d)(w)}] \nsubseteq \Phi_2$ **[NON]**

73. This topic has received next to no attention in recent work on the formal pragmatics of imperatives. Aloni (2007), of course, is interested in permissive content, but fails to connect this to a theory of permissive force. Veltman (2008) is a pleasant exception. Although I believe that our accounts begin from similar intuitions about permissive force, the formalisms do not have anything in common.

74. This actually seems, to me, to require that $c \Vdash P(\phi)(\alpha)$, although I won't actually make this assumption.

(188) If $\Phi \in r_c(d)(w)$ and $\Phi \neq \emptyset$, then, for some $\langle v,v'\rangle \in \Phi$, $\langle v,v'\rangle \in i_c$ **[AV1-R]**

(189) If $\Phi \in r_c(d)(w)$, then $\{v \mid \exists v' : \langle v,v'\rangle \in i_c\} \subseteq \{v \mid \exists v' : \langle v,v'\rangle \in \Phi\}$ **[KE2-R]**

(190) If $\Phi_1 \in r_c(d)(w)$, then if $\Phi_1 \subseteq \Phi_2$, $\Phi_2 \in r_c(d)(w)$ **[EX-R]**

While NDRLs will play no direct role in the static interpretation of deontic modals, they function as the background against which future imperative utterances or grantings of permission may be tested for acceptability. To put it somewhat grandiosely, the permissive force of imperatives is understood in terms of the grant of a sort of Berlin-ian negative right: a freedom to resist certain further kinds of instructions or constraints on her will. Implementing this idea will require redefining updates for imperative or imperative-containing formulas of the language. We begin with imperative formulas.

5.4.1 Imperative Formulas

As before, the easiest way of defining the updates requires allowing proposition-action pairs to occur on an NDRL with respect to an index of evaluation, with the following interpretation:

(191) $\langle p,\Phi\rangle \in r_c(d)(w) \Leftrightarrow \exists\langle v,v'\rangle \in \mathit{min}[i_c \cap (p \times W), \leq t_c(d)(w)] : \Phi \in r_c(d)(v)$

The presence of $\langle p,\Phi\rangle$ on an NDRL at w means that the NDRL at the initial state of each best-from-w p-initial transition contains Φ. We note the following two facts about permission formulas that follow from this definition. (Note also that their converses do not.)

(192) $\langle [\![\phi]\!]_c, [\![\alpha]\!]_c\rangle \in r_c(a_c)(w) \Rightarrow c,w \vDash P(\phi)(\alpha)$

(193) $\forall w \in \{w \mid \exists w' : \langle w,w'\rangle \in i_c\} : \langle [\![\phi]\!]_c, [\![\alpha]\!]_c\rangle \in r_c(a_c)(w) \Rightarrow c \Vdash P(\phi)(\alpha)$

Because we are explicitly representing rights, we resume assuming that a context c contains a parameter that defines Rights List functions—functions from worlds to sets of actions—for individuals. Substantively, the changes to updates are twofold. First, we will complicate the definedness constraints for updates—in addition to requiring that the updated To-Do Lists be admissible, we will require that (i) the permissions expressed by the imperative not conflict with any prior obligations and (ii) updated NDRLs be admissible. Undefinedness is, speaking roughly, the dynamic reflex of the inability to incorporate an imperative without revising prior constraints and obligations. Second, we will have imperatives updating the Rights List function of the subject of the imperative. Let $c = \langle s_c, a_c, i_c, t_c, r_c\rangle$ be a context. Update potentials for basic imperative formulas $!(\phi)(\alpha)$ whose salient alternatives, $\mathit{alt}(!(\phi)(\alpha))$, are given by $\boldsymbol{A}_1;...;\boldsymbol{A}_n$, are given as follows. (If the reader is hazy on alternatives and the $[\wp]$ permission operator, see §4.11.1.)

(194) $c\|!(\phi)(\alpha)\|$ is defined iff

- $t_{c^*}(a_c)$ satisfies **CON***, **AV2****, **KE2***
- $c \not\Vdash \neg[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n]$
- $r_{c^*}(a_c)$ satisfies **CR**, **NON**, **AV1-R**, **KE2-R**, **EX-R**

(195) If defined, $c\|!(\phi)(\alpha)\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*}, r_{c^*}\rangle$

- s_{c^*}, a_{c^*}, i_{c^*}, and t_{c^*} as defined in §5.3.2
- $r_{c^*}(d) = r_c(d)$, for all $d \neq a_c$
- $r_{c^*}(a_c)$ is defined stepwise. Select $w \in \{w \mid \exists w' : \langle w, w'\rangle \in i_c\}$.
  We execute the following changes to $r_c(d)$.

  - If $\alpha \in \boldsymbol{A}_1$, let $r_{c^*}(a_c)(w) = r_c(a_c)(w) \cup \{\langle [\![\phi]\!]_c, [\![\alpha]\!]_c\rangle\}$
  - If $\alpha \in \boldsymbol{A}_2$ and $\langle v, v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_c \times W) \cap [\![\bigoplus \boldsymbol{A}_1]\!]_c, \leq_{t_{c^*}(a_c)(w)}]$,
    let $r_{c^*}(a_c)(v') = r_c(a_c)(v') \cup \{\langle W, [\![\alpha]\!]_c\rangle\}$

    $\vdots$
  - If $\alpha \in \boldsymbol{A}_n$, $\langle v, v'\rangle \in \mathit{min}[i_c \cap ([\![\phi]\!]_c \times W) \cap [\![\bigoplus \boldsymbol{A}_1; \ldots; \bigoplus \boldsymbol{A}_{n-1}]\!]_c, \leq_{t_{c^*}(a_c)(w)}]$,
    let $r_{c^*}(a_c)(v') = r_c(a_c)(v') \cup \{\langle W, [\![\alpha]\!]_c\rangle\}$

  Rinse, later, repeat, for all $w \in \{w \mid \exists w' : \langle w, w'\rangle \in i_c\}$.

This is a complicated definition, but only two things are really going on. First, To-Do Lists are updated exactly as before. Second, by defining the update to the NDRL as we do (as well as by enforcing that updated NDRLs satisfy **CR**, **NON**, **AV1-R**, **KE2-R**, and **EX-R**) we ensure that $c^* \Vdash [\wp][\phi][\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n]$, whenever update is defined. The proof of this fact is omitted, but is obvious enough from the semantics for $[\wp][\phi][\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n]$ together with (192) and (193). Three things to note about NDRL update:

i. (170) is immediate: if defined, $c\|!(\phi)(\alpha) \Vdash \nabla(\phi)(\mathit{alt}(!(\phi)(\alpha)))$, so $c\|!(\phi)(\alpha) \Vdash\ !(\phi)(\alpha)$.

ii. Choice-offering "disjunctive" imperatives of arbitrary complexity introducing genuine rights (and rather complicated conditional rights) corresponding to the relevant "disjuncts" (cf. §4.11). This will be handy for avoiding a dynamic version of the Ross Paradox when defining a dynamic analogue of the static entailment relation $\Vdash$.

iii. The granting of rights strengthens the definedness constraints for updating the context: in view of **NON**, only constraints that are consistent with the exercise of preexisting rights at a world may be added to an agent's To-Do List at that world.

In defining updates for more complicated imperative formulas, some shorthand will be useful. If $i \subseteq i_c$, then we let $r_c(a_c) \pm_i [\wp][\phi][\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n]$ be the result of altering $r_c(a_c)$ by executing the above procedure on each $w \in \{w \mid \exists w' : \langle w, w'\rangle \in i\}$ and leaving $r_c(a_c)$ otherwise unchanged.[75] We note, without proving, the following property of NDRL update.

(196) If $[\wp][\phi][\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n] \vDash [\wp][\phi][\boldsymbol{B}_1; \ldots; \boldsymbol{B}_m]$, then
$\Phi \in r_c(a_c) \pm_i [\wp][\phi][\boldsymbol{B}_1; \ldots; \boldsymbol{B}_m](w)$ implies $\Phi \in r_c(a_c) \pm_i [\wp][\phi][\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n](w)$

Informally, (196) follows from the fact that $r_c(a_c) \pm_i [\wp][\phi][\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n](w)$ differs from $r_c(a_c)(w)$ only with respect to the addition of actions that $[\wp][\phi][\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n]$ says is permit-

75. Although we are not concerned with performatives with only permissive interpretations (grants of permission and the like), it is natural to think they could be handled in terms of updating an NDRL with a $[\wp]$ formula using the $\pm$ operation.

ted at w. If $[\wp][\phi][\boldsymbol{B}_1;...;\boldsymbol{B}_m]$ expresses a permission at w, then that permission is a weakening of a permission that $[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n]$ expresses holds at w, since $[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n] \vDash [\wp][\phi][\boldsymbol{B}_1;...;\boldsymbol{B}_m]$. The claim then follows by the closure of NDRLs under arbitrary expansion (**EX-R**).

5.4.2 Basic Dynamic Formulas

We redefine updates for formulas of the form $[\beta](!(\phi)(\alpha))$, where $\textit{alt}(!(\phi)(\alpha)) = \boldsymbol{A}_1;...;\boldsymbol{A}_n$, so that the updated context validates $[\beta]([\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n])$.

(197) $c\|[\beta](!(\phi)(\alpha))\|$ is defined iff
- $t_{c^*}(a_c)$ satisfies **CON***, **AV2****, **KE2***
- $c \nVdash [\beta](\neg[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n])$
- $r_{c^*}(a_c)$ satisfies **CR**, **NON**, **AV1-R**, **KE2-R**, **EX-R**

(198) When defined, $c\|[\beta](!(\phi)(\alpha))\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*}\rangle$
- s_{c^*}, a_{c^*}, i_{c^*}, and t_{c^*} as defined in §5.3.3
- $r_{c^*}(d) = r_c(d)$, for all $d \neq a_c$
- $r_{c^*}(a_c) = r_c(a_c) \pm_{i_\beta} [\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n]$, where $i_\beta = \{v \mid \exists w : \langle w,v\rangle \in i_c \cap [\![\beta]\!]_c\}$

Essentially the same operation on the context is performed in this case as in the case of basic imperative formulas—alteration of the NDRL is simply restricted to states that are accessible from some possible initial state by the execution of β. The following consequences are immediate.

(199) $c\|[\beta](!(\phi)(\alpha))\| \Vdash [\beta]([\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n])$, when $c\|[\beta](!(\phi)(\alpha))\|$ is defined

(200) $c\|[\beta](!(\phi)(\alpha))\| \Vdash [\beta](!(\phi)(\alpha))$, when $c\|[\beta](!(\phi)(\alpha))\|$ is defined

5.4.3 Complex Dynamic Formulas

Finally, we redefine the update potential for $\lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))$, where $\textit{alt}(!(\phi)(\alpha)) = \boldsymbol{A}_1;...;\boldsymbol{A}_n$, so that the updated context validates $\lceil\!\lceil\beta\rceil\!\rceil^\pi([\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n])$.

(201) $c\|\lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))\|$ is defined iff
- $t_{c^*}(a_c)$ satisfies **CON***, **AV2****, **KE2***
- $c \nVdash \lceil\!\lceil\beta\rceil\!\rceil^\pi(\neg[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n])$
- $r_{c^*}(a_c)$ satisfies **CR**, **NON**, **AV1-R**, **KE2-R**, **EX-R**

(202) When defined, $c\|\lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))\| = c^* = \langle s_{c^*}, a_{c^*}, i_{c^*}, t_{c^*}\rangle$
- s_{c^*}, a_{c^*}, i_{c^*}, and t_{c^*} as defined in §5.3.4
- $r_{c^*}(d) = r_c(d)$, for all $d \neq a_c$
- $r_{c^*}(a_c) = r_c(a_c) \pm_{i_{\pi/\beta}} [\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n]$,
  where $i_{\pi/\beta} = \{v' \mid \exists w \exists v : \langle v,v'\rangle \in \textit{min}[i_c \cap ([\![\pi]\!]_c \times W) \cap [\![\beta]\!]_c, \leq_{t_{c^*}(a_c)(w)}]\}$

Once again, the relevant update is restricted, in this case to terminal states of π-initial transitions of β that are best with respect to the updated To-Do List $t_{c^*}(a_c)$, from the vantage of some w that is compatible with c's picture of the current state of the world. The following consequences of this definition are again immediate.

(203) $c\|\lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))\| \Vdash \lceil\!\lceil\beta\rceil\!\rceil^\pi([\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n])$, when $c\|\lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))\|$ is defined

(204) $c\|\lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))\| \Vdash \lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))$, when $c\|\lceil\!\lceil\beta\rceil\!\rceil^\pi(!(\phi)(\alpha))\|$ is defined

5.4.4 Conceptual Virtues Redux

Let's sum up the attractive features of this system. We predict that:

- When updates are defined, updating a context with an imperative formula of any stripe reliably introduces both the proper sorts of obligations and rights—most significantly, precisely the sorts of rights that appear be granted by arbitrarily complex choice-offering imperatives—on their subjects.

- There is a distinction between performative, permissive, and assertive force. We can, moreover, explain precisely what this difference amounts to, by appeal to different pieces of the semantic and pragmatic apparatus. Assertive force consists in update of the modal base. Imperative force consists in update of a To-Do List. Permissive force consists in the granting of negative rights—rights to resist future attempts to constrain a subject's behavior—which we model with the alteration of an NDRL.

Since we have incorporated permissions, we are able to give a fully general formulation of a constraint semantics for the imperative language. This is worth doing for two reasons: (i) it reveals the connection between the static account we spent most of the paper developing and our dynamic account of imperative force, but additionally (and interestingly) (ii) it reveals an irreducibly dynamic aspect of the account which cannot be expressed simply by listing formulas of either the imperative or deontic object languages that must be validated in order for update to be vacuous. We associate an arbitrary imperative or imperative-containing formula π of $\mathcal{L}_{ILA}$ with a function $\chi_c(\pi) : (\mathcal{T} \times \mathcal{T}) \mapsto \{0,1\}$ of a set characterizing the set $\mathsf{inad}_{\chi_c}(\pi)$ of inadmissible To-Do List/NDRL pairs.

(205) $\mathsf{inad}_{\chi_c}(\pi) = \{\langle \tau_{c,a_c}, \rho_{c,a_c} \rangle \mid \tau_{c,a_c} \nVdash \pi \text{ or } \rho_{c\|\pi\|,a_c} \neq \rho_{c,a_c}\}$

A pair $\langle \tau_{c,a_c}, \rho_{c,a_c} \rangle$ is inadmissible in case either τ_{c,a_c} or ρ_{c,a_c} is inadmissible. Update is triggered at a context c just in case either the subject's To-Do List or NDRL is inadmissible. As before, To-Do Lists are inadmissible if and only if they fail to validate π (if and only if $\tau_{c\|\pi\|,a_c} \neq \tau_{c,a_c}$). There is, however, no corresponding static-semantic condition for inadmissible NDRLs: since NDRLs play no direct role in the static semantics, we are unable to state an equivalent formulation of the necessary and sufficient condition on inadmissible NDRLs (that $\rho_{c\|\pi\|,a_c} \neq \rho_{c,a_c}$) in terms of the static semantics. To illustrate: let $\pi = !(\phi)(\alpha)$, and let $alt(\pi) = \boldsymbol{A}_1;...;\boldsymbol{A}_n$. Even when c validates $[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n]$, it is often the case that $r_c(a_c) \pm_i [\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n] \neq r_c(a_c)$. Rights are finer-grained than permissions: while every right is a permission, not every permission is a right. Being granted a right induces changes in parameters of the context that cannot be fully characterized in terms of the validity of certain formulas in either the deontic or imperative object languages. While updating a context with π always (when defined) yields a context that validates $[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n]$, inadmissible NDRLs cannot be characterized in terms of failure or their corresponding context to validate $[\wp][\phi][\boldsymbol{A}_1;...;\boldsymbol{A}_n]$.

In short, incorporating permission into the dynamics means giving up the pleasant generalization with which we concluded §5.3.5—that the dynamic side of our account models diachronic satisfaction of synchronic constraints on contextual parameters that can be characterized using the static semantics. Constraints on To-Do lists can still be glossed in this way, but constraints on NDRLs cannot. I do not, to be sure, want to portray this as any sort of loss. There is still a genuine connection between the static semantics and the theory of imperative force, and I take it that this is a real virtue of the semantics. It turns out, moreover, that there is an interesting dimension of imperative dynamics—grants of permission—that doesn't reduce to imperative statics. But this, I think, is all to the good: the dynamic semantics we give in the following section will be genuinely dynamic, not a mere notational variant of the static semantics.[76]

5.5 *Dynamic Satisfaction and Entailment*

We have defined the dynamic interpretation function $\|\cdot\|$ for every formula of $\mathcal{L}_{ILA}$. This is the most difficult step in defining a dynamic satisfaction/requirement relation—the rest is standard, and can be presented quickly.[77] A formula ϕ of $\mathcal{L}_{ILA}$ is satisfied by a context c—$c \Subset \phi$—just in case c already bears the information that ϕ (if ϕ is non-imperative) or updating with ϕ imposes no new constraints and confers no new rights on a_c (if ϕ is imperative). Just in case, that is to say, updating c with ϕ leaves the context as before.

(206) $c \Subset \phi$ iff $c\|\phi\| = c$

Note that *neither* of the following relationships hold between dynamic $\Subset$ and static $\Vdash$.

(207) $c \Vdash \phi$ implies $c \Subset \phi$

(208) $c \Subset \phi$ implies $c \Vdash \phi$

This is not, of course, surprising in the case of non-imperative formulas of $\mathcal{L}_{ILA}$. Truth of ϕ at c does not imply presupposition that ϕ at c, nor vice versa. Things are different, however, if we restrict our attention to imperative ϕ: we find that (208) holds, although (207) does not. Regarding (207): this might be thought surprising, since if an imperative ϕ already is in force at c—$c \Vdash \phi$—one might expect that updating the context with ϕ would leave the context as before. But we know better: the permissions expressed by ϕ may be in force at c, yet not counts as rights there. Regarding (208): this is an immediate consequence of our general result that $c\|\phi\| \Vdash \phi$, when ϕ is imperative and $c\|\phi\|$ is defined.

Dynamic entailment in $\mathcal{L}_{ILA}$ is also standard. A list $\phi_1, ..., \phi_n$ of formulas entails ψ—$\phi_1, ..., \phi_n \Subset \psi$—just in case updating a context with $\phi_1, ..., \phi_n$ always (when defined) yields a context that satisfies/requires ψ.

(209) $\phi_1, ..., \phi_n \Subset \psi$ iff, for all c: $c\|\phi_1\|...\|\phi_n\|$ is defined implies $c\|\phi_1\|...\|\phi_n\| \Subset \psi$

Note the following similarity between the dynamic and static entailment relations: an imperative formula entails another just in case the former's command content and permissive content are each at least strong as the latter's. Consider any two imperative formulas $!(\phi)(\alpha)$ and $!(\phi)(\beta)$ such that $alt(!(\phi)(\alpha)) = \boldsymbol{A}_1; ...; \boldsymbol{A}_n$ and $alt(!(\phi)(\beta)) = \boldsymbol{B}_1; ...; \boldsymbol{B}_n$. The claim receives the following formal statement.

(210) $!(\phi)(\alpha) \Subset !(\phi)(\beta)$ iff $O(\phi)(\alpha) \vDash O(\phi)(\beta)$ and $[\wp][\phi][\boldsymbol{A}_1; ...; \boldsymbol{A}_n] \vDash [\wp][\phi][\boldsymbol{B}_1; ...; \boldsymbol{B}_n]$

Proof. We suppose that the relevant updates are all defined. Let c be any context.

$\Rightarrow$: Suppose $!(\phi)(\alpha) \Subset !(\phi)(\beta)$.

Then $t_{c\|!(\phi)(\alpha)\|\|!(\phi)(\beta)\|} = t_{c\|!(\phi)(\alpha)\|}$, and $r_{c\|!(\phi)(\alpha)\|\|!(\phi)(\beta)\|} = r_{c\|!(\phi)(\alpha)\|}$.

Then $\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle \in t_c(a_c)(w)$ implies $\langle [\![\phi]\!]_c, [\![\beta]\!]_c \rangle \in t_c(a_c)(w)$.

76. The distinction between reducibly and irreducibly dynamic accounts is taken from von Fintel & Gillies (2007).

77. Classic references for this sort of approach are Groenendijk & Stokhof (1991); Veltman (1996).

Suppose $c, w \vDash O(\phi)(\alpha)$. Then, by (174), $\langle [\![\phi]\!]_c, [\![\alpha]\!]_c \rangle \in t_c(a_c)(w)$.

So $\langle [\![\phi]\!]_c, [\![\beta]\!]_c \rangle \in t_c(a_c)(w)$.

Then, by (174), $c, w \vDash O(\phi)(\beta)$.

A similar argument shows that $[\wp][\phi][\boldsymbol{A}_1; ...; \boldsymbol{A}_n] \vDash [\wp][\phi][\boldsymbol{B}_1; ...; \boldsymbol{B}_n]$.

$\Leftarrow$: Let $O(\phi)(\alpha) \vDash O(\phi)(\beta)$ and $[\wp][\phi][\boldsymbol{A}_1; ...; \boldsymbol{A}_n] \vDash [\wp][\phi][\boldsymbol{B}_1; ...; \boldsymbol{B}_n]$.

To show: $t_{c\|!(\phi)(\alpha)\|\|!(\phi)(\beta)\|} = t_{c\|!(\phi)(\alpha)\|}$ and $r_{c\|!(\phi)(\alpha)\|\|!(\phi)(\beta)\|} = r_{c\|!(\phi)(\alpha)\|}$.

The latter follows immediately from (196).

As for the former, since $c\|!(\phi)(\alpha)\| \Vdash O(\phi)(\alpha)$, $c\|!(\phi)(\alpha)\| \Vdash O(\phi)(\beta)$.

So, by (175) and the update potential for $!(\phi)(\beta)$, $t_{c\|!(\phi)(\alpha)\|\|!(\phi)(\beta)\|} = t_{c\|!(\phi)(\alpha)\|}$.

Note that we have an analogue of this result for the static imperative entailment relation $\Vvdash$, which follows from the analysis of static permissive content presented in §4.11.1.

(211) $!(\phi)(\alpha) \Vvdash !(\phi)(\beta)$ iff $O(\phi)(\alpha) \vDash O(\phi)(\beta)$ and $[\wp][\phi][\boldsymbol{A}_1; ...; \boldsymbol{A}_n] \vDash [\wp][\phi][\boldsymbol{B}_1; ...; \boldsymbol{B}_n]$

This is, in the static context, the fact we have exploited in order to address the Ross Paradox (and arbitrarily complicated versions thereof). The invalidity of the Ross inference is explained by the fact that, e.g., $!(\phi)(\alpha + \beta)$ generally has richer permissive content than $!(\phi)(\alpha)$. For one imperative to entail another, the permissive content of the former must be at least as strong as the permissive content of the latter. The same explanation is what accounts for dynamic invalidity of the Ross inference, in view of the above result.

So, while dynamic satisfaction of imperative formulas at a context does not reduce to static validity at a context (in view of the falsity of (207) when ϕ is imperative), we nevertheless find that dynamic *entailment* relations between imperative formulas never fail to coincide with static entailment relations between them:

(212) $!(\phi)(\alpha) \Subset !(\phi)(\beta)$ iff
$!(\phi)(\alpha) \Vvdash !(\phi)(\beta)$ iff
$c \Vdash !(\phi)(\alpha)$ implies $c \Vdash !(\phi)(\beta)$, for arbitrary c

More generally, let π be an arbitrary formula of $\mathcal{L}_{ILA}$ embedding $!(\phi)(\alpha)$—a formula of of either the form $[\beta](...!(\phi)(\alpha)...)$ or $[\![\beta]\!]^{\psi}(...!(\phi)(\alpha)...)$—and let π' be the result of uniformly substituting occurrences of $!(\phi)(\alpha)$ in π with $!(\phi)(\beta)$. It follows from the equivalence of dynamic and static entailment relations between imperative formulas that $\pi \Subset \pi'$ iff $\pi \Vvdash \pi'$.

We cannot, however, fully generalize the equivalence. That is to say, the following does not generally hold for arbitrary formulas of $\mathcal{L}_{ILA}$.

(213) $\phi_1, ..., \phi_n \Subset \psi$ iff $\phi_1, ..., \phi_n \Vvdash \psi$

The equivalence does hold when all of $\phi_1, ..., \phi_n, \psi$ contain an imperative, as follows from the generalization of (212) to arbitrary $\mathcal{L}_{ILA}$ embeddings of $\mathcal{L}_{ILA}$ imperatives. But it fails to hold when a premise is non-imperative. We illustrate the point with a simple example:

(214) $!(\phi)(\alpha), \phi \Subset !(\top)(\alpha)$, while $!(\phi)(\alpha), \phi \not\Vvdash !(\top)(\alpha)$

Proof that $!(\phi)(\alpha), \phi \Subset !(\top)(\alpha)$.

Let $c^* = c\|!(\phi)(\alpha)\|\|\phi\|$, and suppose that ϕ is non-imperative.

We know *min*$[i_{c^*} \cap ([\![\phi]\!]_{c^*} \times W), \leq_{t_{c^*}(a_{c^*})(w)}] \subseteq [\![\alpha]\!]_{c^*}$ and $i_{c^*} \subseteq [\![\phi]\!]_{c^*} \times W$.

So it follows that $i_{c^*} \cap ([\![\phi]\!]_{c^*} \times W) = i_{c^*}$. Then *min*$[i_{c^*}, \leq_{t_{c^*}(a_{c^*})(w)}] \subseteq [\![\alpha]\!]_{c^*}$.

So $c^* \Vdash !(\top)(\alpha)$.

By contrast, we only have $!(\phi)(\alpha), \phi \Vvdash !(\top)(\alpha)$ if we "dynamicize" the static entailment relation, so that information to the turnstile's left is added to the modal base against which the truth or requirement of the formula to the turnstile's right is checked (roughly along the lines of the dynamicization suggested toward the end of §3.2). As we noted in §2.4, validating this argument form is extremely natural for an update semantic logic of planning.[78] It is less natural, but still potentially reasonable—depending on whether and when it makes sense to use a dynamicized version of the static entailment relation—from the point of view of a logic of content.

The reason validating this argument form is extremely natural for an update semantic logic of planning points to a sense in which the update semantics given in this section probably cannot be given a logic of content rationale. Update semantic accounts define entailment in terms of vacuous update of an *updater's* cognitive state (or contextual representation thereof). From the point of view of the updater—the subject of authority, the owner of the updated To-Do List—updating with the information that ϕ constrains the relevant possibilities for the current state of the world to ϕ-possibilities. Subsequent update with the imperative $!(\phi)(\alpha)$ requires the subject to, in effect, add the action α to her To-Do List at all the best ϕ-worlds. Subsequent to both updates: all the relevant possibilities for the current state of the world are ϕ-possibilities, and the best ones of those demand α. In view of her knowledge of present circumstances, compliance requires that the subject perform α. Subjects are constrained to obey in view of their own information state. But authorities are certainly not committed to prefer obedience in view of their *subjects'* information state—though perhaps committed to prefer obedience in view of her own, the possibility that the dynamicized entailment relation in §2.4 is designed to accommodate. While updating a subject's cognitive state with $!(\phi)(\alpha)$ and ϕ may leave her with no choice but to perform α, this is something an authority is not generally committed to endorsing (if, suppose, her information should differ from that of her subject). I see no real way, then, to give the subject-oriented update semantics developed in this section a logic of content rationale.

5.6 Conclusion: For Pluralism in Imperative Logic

Theories of performative force, imperative or otherwise, are addressee-oriented: their subject-matter is the effect of the use of some linguistic device *on an addressee*. Because of this, building an update semantics around a theory of imperative performative force means, as we just saw, difficulties in accounting for the sense in which imperative logic is supposed to be normative for an authority's endorsement. Requiring that a semantics for imperatives be an outgrowth of a dynamic theory of imperative performative force, à la Portner (2008), then, appears to mean automatically marginalizing a prima facie reasonable, and certainly worthy of study, approach to the logic of imperatives—that which is characteristic of logics of content.

78. This argument is a rough $\mathcal{L}_{ILA}$ analogue of argument forms (12), (13), and (17). There are less rough analogues, but making the relevant points would involve introducing some complexities about representing future tense indicatives in $\mathcal{L}_{ILA}$ that aren't worth the trouble.

More generally, I think that privileging dynamic semantic analyses of the imperative over static treatments, along the lines developed in the major sections of this paper or otherwise, involves marginalizing a collection of prima facie reasonable approaches to the logic of imperatives. I see no compelling reason for this bias.[79] Static logics of content attempt to model the endorsements that an authority is committed to make at a given point in time, while static logics of planning attempt to model the constraints on a subject's planning that are in force at a given point in time. For a wide array purposes, these things certainly seem like they would be well worth modeling, and modeling them certainly seems to require embracing a static perspective on imperative semantics. Even insofar as we think of imperative performative force as a (perhaps the) crucial aspect of the imperative—as something that any semantic treatment of the imperative should say something about—we have seen, at some length, that the static treatments of the imperative on offer *do* have quite a lot to say on this subject, from the vantage of a constraint-semantic orientation. Even where it proved impossible to characterize inadmissible NDRL pairs by appeal to antecedent static validity of the relevant permission formulas, it is easy to see that NDRL update is designed, in part, to enforce their posterior validity.

In short, a pluralistic attitude toward the semantics of imperatives is warranted. Marginalizing reasonable approaches is unwarranted and, moreover, methodologically suspect, in particular for philosophers and linguists whose primary interest lies in the formal semantics of natural language imperatives. This is an area of research to which relatively little attention has been paid or progress made since the 1960s (although there are, of course, a few notable exceptions), and into which contemporary methods in formal semantics are only beginning to penetrate. We ought to save the worrying about who should be working on what, at least until the number of recent papers in major journals outnumbers the fingers on one hand.

References

Aloni, Maria. 2007. Free choice, modals, and imperatives. *Natural Language Semantics* 15: 65–94. doi:10.1007/s11050-007-9010-2.

Aloni, Maria & Robert van Rooy. 2002. The dynamics of questions and focus. In B. Jackson (ed.) *Proceedings of SALT 12*. CLC Publications. URL http://staff.science.uva.nl/~maloni/salt02/salt02.pdf.

Asher, Nicholas & Alex Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.

Bar-Hillel, Yehoshua. 1966. Imperative inference. *Analysis* 26: 79–82. doi:10.2307/3326286.

Bratman, Michael E. 1987. *Intention, Plans, and Practical Reason*. Cambridge: Harvard University Press.

Broome, John. 1999. Normative requirements. *Ratio* 12: 398–419. doi:10.1111/1467-9329.00101.

Castaneda, Hector Neri. 1958. Imperatives and deontic logic. *Analysis* 19: 42–48. doi:10.2307/3326749.

Castaneda, Hector Neri. 1960. Imperative reasonings. *Philosophy and Phenomenological Research* 21: 21–49. doi:10.2307/2104787.

Castaneda, Hector Neri. 1971. There are command sh-inferences. *Analysis* 32: 13–19. doi:10.2307/3327276.

Charlow, Nathan. 2008a. Imperative semantics and dynamics. Unpublished ms., University of Michigan.

79. I don't believe Portner has furnished one, either in published work or personal correspondence (see, e.g., his contribution to the discussion at http://crapulae.wordpress.com/2009/01/30/portners-mistake-about-imperatives/).

Charlow, Nathan. 2008b. Imperatives in context. Unpublished ms., University of Michigan.

Charlow, Simon. 2008c. Free and bound pro-verbs: A unified treatment of anaphora. In T. Friedman & S. Ito (eds.) *Proceedings of SALT 18*. CLC Publications. URL http://semanticsarchive.net/Archive/mI5YTEyO/.

Chellas, Brian F. 1980. *Modal Logic: An Introduction*. Cambridge: Cambridge University Press.

von Fintel, Kai & Anthony S. Gillies. 2007. An opinionated guide to epistemic modality. In T. Gendler & J. Hawthorne (eds.) *Oxford Studies in Epistemology: Volume 2*. Oxford University Press. URL http://mit.edu/fintel/fintel-gillies-2007-ose2.pdf.

von Fintel, Kai & Sabine Iatridou. 2008. What to do if you want to go to harlem: Anankastic conditionals and related matters. URL http://mit.edu/fintel/www/harlem-rutgers.pdf. Unpublished ms., MIT.

Geach, P. T. 1958. Imperative and deontic logic. *Analysis* 18: 49–56. doi:10.2307/3326785.

Geach, P. T. 1963. Imperative inference. *Analysis* 23, Supplement 1: 36–42. doi:10.2307/3326619.

Geach, P. T. 1966. Dr. Kenny on practical inference. *Analysis* 26: 76–79. doi:10.2307/3326285.

Geurts, Bart. 2005. Entertaining alternatives: Disjunctions as modals. *Natural Language Semantics* 14: 383–410. doi:10.1007/s11050-005-2052-4.

Gillies, Anthony S. 2007. Counterfactual scorekeeping. *Linguistics and Philosophy* 30: 329–360. doi:10.1007/s10988-007-9018-6.

Gillies, Anthony S. 2008. Iffiness. URL http://www-personal.umich.edu/~thony/iffiness-revised-june2008.pdf. Unpublished ms., University of Michigan.

Green, Mitchell. 1997. The logic of imperatives. In E. Craig (ed.) *The Routledge Encyclopedia of Philosophy*, 717–721. New York: Routledge.

Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.

Groenendijk, Jeroen & Martin Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14: 39–100. doi:10.1007/BF00628304.

Hamblin, Charles. 1987. *Imperatives*. Oxford: Basil Blackwell.

Hare, R. M. 1949. Imperative sentences. *Mind* 58: 21–39. doi:10.1093/mind/LVIII.229.21.

Hare, R. M. 1952. *The Language of Morals*. Oxford: Oxford University Press.

Hare, R. M. 1967. Some alleged differences between imperatives and indicatives. *Mind* 76: 309–326. doi:10.1093/mind/LXXVI.303.309.

Isaacs, James & Christopher Potts. 2003. Hidden imperatives. URL http://people.umass.edu/potts/talks/isaacs-potts-nels34-handout-1up.pdf. Talk delivered at NELS 34, SUNY Stony Brook.

Jørgensen, Jørgen. 1937-8. Imperatives and logic. *Erkenntnis* 7: 288–296.

Kenny, A. J. 1966. Practical inference. *Analysis* 26: 65–75. doi:10.2307/3326284.

Kolodny, Niko & John MacFarlane. 2008. Ifs and oughts. URL http://johnmacfarlane.net/ifs-and-oughts.pdf. Unpublished ms., UC Berkeley.

Kratzer, Angelika. 1981. The notional category of modality. In H. Eikmeyer & H. Rieser (eds.) *Words, Worlds, and Contexts*, 38–74. Berlin: De Gruyter.

Kratzer, Angelika. 1991. Conditionals. In A. von Stechow & D. Wunderlich (eds.) *Semantics: An International Handbook of Contemporary Research*, 651–656. Berlin: De Gruyter.

Kratzer, Angelika & Junko Shimoyama. 2002. Indeterminate pronouns: the view from Japanese. In Y. Otsu (ed.) *Proceedings of the Third Tokyo Conference on Psycholinguistics (TCP 2002)*, 1–25. Tokyo: Hituzi Syobo.

Lemmon, E. J. 1965. Deontic logic and the logic of imperatives. *Logique et Analyse* 8: 39–71.

Lewis, David. 1973. *Counterfactuals*. Malden: Basil Blackwell.

Lewis, David. 1975. Adverbs of quantification. In E. Keenan (ed.) *Formal Semantics of Natural Language*, 3–15. Cambridge: Cambridge University Press.

Mastop, Rosja. 2005. What can you do? Ph.D. Dissertation, ILLC.

Ninan, Dilip. 2005. Two puzzles about deontic necessity. In J. Gajewski, V. Hacquard, B. Nickel & S. Yalcin (eds.) *New Work on Modality*, Vol. 51 of *MIT Working Papers in Linguistics*. Cambridge: MITWPL.

Portner, Paul. 2004. The semantics of imperatives within a theory of clause types. In K. Watanabe & R. Young (eds.) *Proceedings of SALT 14*. CLC Publications. URL http://semanticsarchive.net/Archive/mJlZGQ4N/.

Portner, Paul. 2008. Imperatives and modals. *Natural Language Semantics* 15: 351–383.

doi:10.1007/s11050-007-9022-y.

Potts, Christopher. 2003. Keeping world and will apart: A discourse-based semantics for imperatives. URL http://people.umass.edu/potts/talks/potts-nyu-handout.pdf. Talk delivered at NYU Syntax/Semantics Lecture Series.

Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1: 75–116. doi:10.1007/BF02342617.

Ross, Alf. 1941. Imperatives and logic. *Theoria* 7: 53–71.

Schulz, Katrin. 2003. You may read it now or later: A case study on the paradox of free choice permission. URL http://www.illc.uva.nl/Publications/ResearchReports/MoL-2004-01.text.pdf. M.A. Thesis, University of Amsterdam.

Schulz, Katrin. 2005. A pragmatic solution for the paradox of free choice permission. *Synthese* 147: 343–377. doi:10.1007/s11229-005-1353-y.

Schwager, Magdalena. 2006. Conditionalized imperatives. In M. Gibson & J. Howell (eds.) *Proceedings of SALT 16*. CLC Publications. URL http://user.uni-frankfurt.de/~scheiner/papers/schwagerFEB07.pdf.

Segerberg, Krister. 1990. Validity and satisfaction in imperative logic. *Notre Dame Journal of Formal Logic* 31: 203–221. doi:10.1305/ndjfl/1093635415.

Segerberg, Krister. 1994. A model existence theorem in infinitary propositional modal logic. *Journal of Philosophical Logic* 23: 337–367. doi:10.1007/BF01048686.

Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.) *Syntax and Semantics 9: Pragmatics*. New York: Academic Press.

Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 25: 701–721. doi:10.1023/A:1020867916902.

Swanson, Eric. 2006. Interactions with context. Ph.D. Dissertation, MIT.

Swanson, Eric. 2008a. Constraint semantics and its application to conditionals. URL http://www-personal.umich.edu/~ericsw/research/Swanson,%20Constraint%20Semantics%20handout.pdf. Talk delivered at First Formal Epistemology Festival, University of Konstanz.

Swanson, Eric. 2008b. Modality in language. *Philosophy Compass* 3: 1193–1207. doi:10.1111/j.1747-9991.2008.00177.x.

Veltman, Frank. 1996. Defaults in update semantics. *Journal of Philosophical Logic* 25: 221–261. doi:10.1007/BF00248150.

Veltman, Frank. 2008. Imperatives at the semantics/pragmatics interface. Unpublished ms., ILLC.

Vorobej, Mark. 1982. Deontic accessibility. *Philosophical Studies* 41: 317–319. doi:10.1007/BF00353882.

Williams, B. A. O. 1963. Imperative inference. *Analysis* 23, Supplement 1: 30–36. doi:10.2307/3326619.

Yalcin, Seth. 2007. Epistemic modals. *Mind* 116: 983–1026. doi:10.1093/mind/fzm983.

Zimmermann, Thomas Ede. 2000. Free choice disjunction and epistemic possibility. *Natural Language Semantics* 8: 255–90. doi:10.1023/A:1011255819284.