J Philos Logic
DOI 10.1007/s10992-013-9284-4

Logic and Semantics for Imperatives

Nate Charlow

Received: 11 May 2012 / Accepted: 26 March 2013

Abstract In this paper I will develop a view about the semantics of imperatives, which I term Modal Noncognitivism, on which imperatives might be said to *have* truth conditions (dispositionally, anyway), but on which it does not make sense to see them as *expressing propositions* (hence does not make sense to *ascribe to them* truth or falsity). This view stands against "Cognitivist" accounts of the semantics of imperatives, on which imperatives are claimed to express propositions, which are then enlisted in explanations of the relevant logico-semantic phenomena. It also stands against the major competitors to Cognitivist accounts—all of which are non-truth-conditional and, as a result, fail to provide satisfying explanations of the fundamental semantic characteristics of imperatives (or so I argue). The view of imperatives I defend here improves on various treatments of imperatives on the market in giving an *empirically* and *theoretically* adequate account of their semantics and logic. It yields explanations of a wide range of semantic and logical phenomena about imperatives—explanations that are, I argue, at least as satisfying as the sorts of explanations of semantic and logical phenomena familiar from truth-conditional semantics. But it accomplishes this while defending the notion—which is, I argue, substantially correct—that imperatives could not have propositions, or truth conditions, as their *meanings*.

This work has benefited greatly from the input of others. I would like to thank Andrew Alwood, Chris Barker, Simon Charlow, Allan Gibbard, Benj Hellie, Paul Portner, Will Starr, Eric Swanson, and Richmond Thomason, as well as audiences at the Association for Symbolic Logic's 2011 APA Group Session on Dynamic Semantics, Cornell University, École Normale Supérieure, NYU, University of Toronto, and York University. Thanks also to two anonymous referees for this journal.

N. Charlow (✉)
Department of Philosophy, University of Toronto,
170 St. George Street, Toronto, ON M5R 2M8, Canada
e-mail: nate.charlow@gmail.com

Published online: 16 May 2013

Keywords Imperatives · Semantics of imperatives · Imperative logic · Deontic logic · Conditional imperatives · Noncognitivism · Expressivism

1 Introduction

After something of a heyday in the third quarter of the Twentieth Century,[1] work on the semantics and logic of imperatives—for our purposes, a sentence like (1) of the form $!\phi$—is presently enjoying a renaissance in both philosophy[2] and linguistics.[3]

(1) Shut the window!

Much of this interest is presumably due to the fact that, while there is a well-established paradigm—the truth-conditional, or model-theoretic, paradigm—for theorizing about the meaning of informational discourse (especially declaratives, but also interrogatives), there is as yet *no such paradigm for purely practical discourse* (imperatives), in either philosophy or linguistics. While there are some nascent ideas about how to begin theorizing about imperatives—e.g., Cognitivist approaches, which rather unambitiously (and, I will suggest, implausibly) treat the semantic values of imperatives as ordinary propositions, and the familiar Noncognitivist, Speech Act-Theoretic, and Expressivist approaches, which rather immodestly (and also, I will suggest, implausibly) ask us to rewrite much of the practice of semantic theorizing from the ground up—all ultimately seem unattractive when scrutinized.

To this end, this paper develops a novel view about the semantics of imperatives (built on top of a relatively new way of understanding the relationship that obtains between a sentence's semantics and its cognitive function). On my view, imperatives might be said to *have* truth conditions (in roughly the way that predicates have truth conditions—more specifically, truth-*of-an-object conditions*). It will not, however, make sense to *ascribe to them* truth or falsity (in the way that it does not make sense to call a predicate true or false *simpliciter*). The meaning of an imperative of the form $!\phi$ is identified with a *property of a plan*: the property it has just if it requires or is settled on seeing to it that ϕ. This account of imperative meaning yields, I argue, a satisfactory account of imperatives' semantic and logical characteristics—one at least as satisfactory as the best truth-conditional account. It does so while avoiding well-known philosophical difficulties that arise from identifying the meaning of an imperative with a proposition or truth condition. And it offers hope of grounding a fruitful research program into the semantics of the wide variety of imperative constructions (conditional and otherwise embedded imperatives) that we observe in natural language.

[1] Exemplary references include [12, 17, 24, 32–34, 41, 48, 50, 62, 87].

[2] See, e.g., [57, 75, 84–86]. In addition to this, there is a great deal of work on imperatives and similar topics that originates from meta-ethical concerns; a small sampling: [3, 18, 37, 49, 65].

[3] See, e.g., [2, 5, 40, 47, 56, 58, 59, 68, 83].

Here, briefly, is an overview of this paper's structure. In Section 2, I describe a class of core phenomena about imperatives (and related kinds of sentences, e.g., permissions) having to do with semantic properties (inconsistency), relations (logical consequence), and computations (compositional derivation of complex meanings) that apply to sentences or inferences that involve imperatives. I go on to review why the most obvious non-truth-conditional approaches to these phenomena are lacking, in ways that have everything to do with the fact that they are not truth-conditional in nature.

In Section 3, I look at two ways of assigning imperative truth conditions, each of which qualifies as a form of Cognitivism. One of these approaches—the approach that assigns imperatives *modal* logical forms—seems to do *very well* with respect to the phenomena identified in Section 2.

In Section 4, I explain why the apparent success of Modal Cognitivism is, ultimately, an illusion. The explanations Modal Cognitivism claims to offer, while they possess a certain formal allure, are, in fact, impossible to state intelligibly in the semantic metalanguage (and, so, do not qualify as *explanations* of anything at all). Modal Cognitivism (and Cognitivism, generally) also runs afoul of the platitude that a sentence's semantic value should explain why that sentence is suited for the cognitive and discursive roles that it actually has.

In Section 5, I develop a new approach to theorizing about the semantics of imperatives: what I call *Modal Noncognitivism*. Modal Noncognitivism retains the apparent virtues of Modal Cognitivism—indeed, transforms these apparent virtues into actual ones—while avoiding the problems that arise from treating the semantic value of an imperative as a proposition. Modal Noncognitivism is rooted in a new approach to semantic theorizing, due especially to [77, 78], on which sentences are, in general, held to express (or otherwise characterize) a *cognitive instruction* for an agent. While the fruits of this approach are largely uninteresting for descriptive (propositional) language, they are absolutely central to the design of an adequate semantics for imperatives.

Imperatives, I claim, express a semantic value uniquely suited to resolving a cognitive issue for an agent: the issue of *what to plan* or *how to act*. While declaratives, in virtue of expressing propositions, encode what we might term a *locational perspective* for an agent, imperatives encode what we might term a *directional* perspective. In the case of basic imperatives, that semantic value is, unsurprisingly, just the property that a plan instantiates when it is decided on realizing the thing recommended by the imperative. This little idea bears considerable fruit when it comes to theorizing about the semantics of things, like permissions, that seem intimately semantically related to imperatives, but which are resistant to the techniques of ordinary forms of Noncognitivism (e.g., those which assign imperatives semantic values which permissions manifestly lack, like compliance-conditions). Permissions, like imperatives, semantically characterize properties of plans—namely, the property of permitting the thing which the permission seems to permit. Their relation to imperatives is a logical relation, in view of the fact that, just as a matter of logic, plans that require things permit certain things, and fail to permit others.

Finally, I show how idea also extends naturally to treatments of *embedded imperatives* (most notably, *conditional imperatives*), which have proved most resistant to

semantic analysis and understanding. Modal Noncognitivism yields a plausible, compositional account of the meaning of conditional imperatives, while also suggesting a general strategy for computing the meaning of any complex construction embedding an imperative.

Modal Noncognitivism is, I conclude, a powerful theory for imperatives—one on which future work in semantics, logic, and philosophy can (and should) build.

2 The Utility of Truth Conditions

Truth conditions are useful for accounting for various sorts of semantic properties (e.g., consistency), relations (e.g., consequence), and operations (e.g., compositional computation of meaning). For instance, given a suitable notion of compatibility (e.g., co-satisfiability in a model or valuation), the following notions are standard:

Truth conditional consistency
ϕ and ψ are consistent iff ϕ's truth condition is compatible with ψ's truth condition

Truth conditional consequence
ψ is a consequence of $\phi_1, ..., \phi_n$ iff the truth condition of $\neg\psi$ is incompatible with the truth condition of $\phi_1 \wedge ... \wedge \phi_n$

The following understanding of compositional semantic computation is also familiar.

Truth conditional compositionality
If ϕ is a sentence, ϕ's semantic contribution to any environment in which it is embedded is its truth condition

The utility of the truth conditional paradigm for theorizing about (in)consistency and consequence for propositional and first-order languages, as well as for compositionally computing the meanings of complex syntactic constructions, is well-established. It is also well-established for theorizing about such semantics notions for languages that are built "on top of" such languages, like, e.g., modal languages.

Like modal languages, imperative languages, with some exceptions that we will for the most part ignore here,[4] are usually thought to be built "on top of" truth-conditional "base" languages (e.g., propositional or first order languages). For instance, the following might constitute a definition of the most basic "mixed" imperative language.

Let $\mathcal{L}$ be a first-order language. Then the imperative language $\mathcal{L}_{\texttt{imp}}$ is the smallest set such that:

(i) $\mathcal{L} \subseteq \mathcal{L}_{\texttt{imp}}$
(ii) If $\phi \in \mathcal{L}$, then $!\phi \in \mathcal{L}_{\texttt{imp}}$.

[4] See, e.g., [6] and subsequent work in the STIT tradition.

Such a language would contain both the language of first-order logic, as well as sentences constructible by prefixing the imperative operator ! to a sentence of first-order logic, but (following the relatively standard dictum that imperatives are representable with operators indicating imperative "force" scoping over sentences of some base language[5]) no sentence embedding an imperative sentence under quantifiers or connectives, and no sentence involving iterated occurrences of !.

Even given such a primitive imperative language, it might already seem a short step to using truth conditional understandings of consistency and consequence to explain the following facts about imperatives.

- **Fact 1.** They stand in *inconsistency relations* with other sentences (imperative and otherwise)
- **Fact 2.** They stand in *inferential relations* with other sentences (imperative and otherwise): an imperative can, in some cases, be validly inferred as the conclusion of an argument

And, in view of the following fact...

- **Fact 3.** They *embed* in fairly regular ways (most clearly, as the consequents of indicative conditionals)

...we might find ourselves—once we have extended $\mathcal{L}_{\texttt{imp}}$ to allow the representation of such sentences—reaching for the familiar idea that sentences contribute *their truth conditions* to the computation of the meaning of complex constructions to account for the meanings of, e.g., indicative conditionals with imperative consequents.

Let me expand on these gestures in order.

2.1 Imperative Inconsistency

To illustrate Fact 1, the following pairs seem to be fairly obviously inconsistent, in the most familiar sense of that notion.

(2) a. Move that rock!
b. Don't move that rock!

(3) a. Go ahead, take an apple! (*permission imperative*; see [61])
b. Don't you dare take an apple!

(4) a. Go to your room!
b. Go to your room or don't (...I don't care) (*choice-offering imperative*; see [4, 61])

What kind of inconsistency? What I mean by this is that it seems clear that these sentences are inconsistent, in ways that *do not differ starkly* from the sort of inconsistency exhibited by sentences of the form ϕ and $\neg\phi$.

- Anyone who issues them can rightly be charged with inconsistency.

[5] For views in this tradition, see [33]; [[88] Section 23]; [22, 50, 76, 92].

– Anyone who accepts them (by which I do *not* mean that they come to *believe that* they are binding, rather, that they try to adjust their desires and plans accordingly) can rightly be charged with inconsistency.
– Any individual who issues one disagrees with an individual who issues the other. The former disagrees with the latter about what the addressee is to do.

It would seem overwhelmingly plausible that such facts about such pairs of imperatives, like the analogous facts about non-imperative sentences and their negations, have a semantic (rather than, say, pragmatic or illocutionary) origin. Although they certainly give rise to pragmatic and illocutionary inconsistencies (when, e.g., a single speaker issues contradictory imperatives), such inconsistencies are special cases—it is doubtful that all of the above forms of inconsistency (e.g., the sort of inconsistency afflicting someone who accepts conflicting imperatives) could be explained in terms of them. Instead, we should attempt to explain these special kinds of inconsistency as having a fundamentally semantic origin (as, for instance, we typically explain the inconsistency of an agent who accepts sentences of the form ϕ and $\neg\phi$ being grounded in the semantic inconsistency of a sentence and its negation).

Permissions, too Suppose, inspired perhaps by [53], we find it natural to extend the imperative language $\mathcal{L}_{\texttt{imp}}$ with a clause for dedicated permission sentences, as follows:

Let $\mathcal{L}$ be a first-order language. Then the imperative language $\mathcal{L}^{+}_{\texttt{imp}}$ is the smallest set such that:

(i) $\mathcal{L} \subseteq \mathcal{L}^{+}_{\texttt{imp}}$
(ii) If $\phi \in \mathcal{L}$, then $!\phi \in \mathcal{L}^{+}_{\texttt{imp}}$ and $¡\phi \in \mathcal{L}^{+}_{\texttt{imp}}$.

Sentences of the form $¡\phi$ have the same sort of meaning that is assigned to permission imperatives like (3a) and performative (i.e., permission-generating) interpretations of sentences like (5) and (6).

(5) You may have an apple
(6) I permit you to have an apple

English, of course, despite (like, as far as we know, all natural languages) having a dedicated clause-type for the expression of positive desires concerning the actions of one's addressee (i.e., the imperative), lacks a dedicated clause-type for the expression of permissions.[6] But clearly a formal language with such a dedicated clause-type is both definable and intelligible. And we will expect such a language to witness semantic facts like the following.

– The inconsistency of $!\phi$ and $¡\neg\phi$
– The inconsistency of $!\neg\phi$ and $¡\phi$

[6] On the notion of clause-type and the universality of the imperative, see especially [58].

It seems plausible, for reasons already seen, that such facts would be semantic in origin.

2.2 Imperative Consequence

Fact 2 is more complicated to motivate, but at least this inference seems to be clearly valid:

(7) a. Move all the rocks!
b. So, move that rock!

That might seem to suggest, generally, that an imperative of the form $!\forall x Fx$ entails an imperative of the form $!Fa$. Similarly, an imperative of the form $!(\phi \wedge \psi)$ seems to entail $!\phi$.

(8) a. Take out the trash and mow the lawn!
b. So, take out the trash!

Can we say generally, then, than an imperative of the form $!\phi$ entails every imperative of the form $!\psi$, when ψ is an ordinary consequence of ϕ? This is more complicated.

A note on paradoxes Unsurprisingly, given the obvious connection between obligation and commanding, there are paradoxes of imperative inference that closely resemble certain paradoxes of inference in deontic logic. For instance, there is Ross' [63] Paradox...

(9) a. Post the letter! $\approx !\phi$
b. #So, post or burn the letter! $\approx !(\phi \vee \psi)$
(10) a. You ought to post the letter $\approx O\phi$
b. #So, you ought to post or burn the letter $\approx O(\phi \vee \psi)$

...as well as closely related paradoxes having to do with *closure properties* of both imperative and obligation operators.

(11) a. Put the stunt double in a protective suit and light him on fire! $\approx !(\phi \wedge \psi)$
b. #So, light the stunt double on fire! $\approx !\psi$
(12) a. You ought to put the stunt double in a protective suit and light him on fire $\approx O(\phi \wedge \psi)$
b. #So, you ought to light the stunt double on fire $\approx O\psi$

Here we have two kinds of apparent counterexample to the generalization that an imperative of the form $!\phi$ entails every imperative of the form $!\psi$, when ψ is a consequence of ϕ. These are paralleled by apparent counterexamples to the generalization that a deontic of the form $O\phi$ entails every deontic of the form $O\psi$, where ψ is an ordinary consequence of ϕ.

What is the relevance of these "paradoxes" of imperative and deontic implication for our project here? Since analogous puzzles arise for deontic inferences (and since no one thereby takes this to jeopardize the project of giving a semantic account

of deontic entailment), I do not take the existence of such puzzles for imperative inferences to jeopardize the project of giving a semantic account of imperative consequence. There are valid imperative inferences—for instance, the inference from $!\forall x\, F x$ to $!Fa$ and from $!(\phi \wedge \psi)$ to $!\phi$ (although *not* to $!\psi$ [7])—and imperative validity merits a semantic account.[8]

Such puzzles do *complicate* the project of giving such a semantic account of imperative consequence (just as they complicate the project of giving a semantic account of deontic consequence). Since, however, there is manifestly some sort of semantic phenomenon here, it is manifestly worthwhile to try to lay down some sort of paradigm for theorizing about semantic properties and relations for a language that contains imperatives.

That is what I will attempt to do in the remainder of this paper. Although I will gesture at how such paradoxes might be resolved for imperatives—there is, I will suggest, no special puzzle at all for the semantics of *imperatives*—I will not give the issue any sustained attention. Laying down a semantic paradigm is my main interest. Dealing with paradoxes is a project to be tackled only after such a paradigm is in place.

2.3 Embedding

Contrary to a (mysteriously large) body of received wisdom, there is a great deal of evidence suggesting that imperatives are *irreducibly embeddable* (see [13, 67, 68, 75]). What I mean by "irreducibly embeddable" is that there are imperatives in natural language that cannot be represented with a sentence from the primitive imperative language we defined above, in which the imperative operator takes widest scope in

[7] Since the standard interpretation of $\wedge$ is commutative, this would seem to commit me to some degree of non-substitutivity-of-logical-equivalents principle for imperative contexts. In reality, I (along with a large contingent of working semanticists) think natural language conjunction is *not* in general commutative. Thus, if $\wedge$ is to represent natural language conjunction, $\wedge$ should not be given a commutative interpretation; if it is not to represent natural language conjunction (but, instead, the familiar Boolean connective), it is not generally correct to represent natural language conjunction using $\wedge$. In particular, I think that *and*, when it coordinates distinct descriptions of actions by way of describing a complex action, should be understood as non-commutative: lighting the stunt double on fire and putting him in a protective suit is a *much* different action from putting the stunt double in a protective suit and lighting him on fire. It is, therefore, misleading to represent an imperative like *put the stunt double in a protective suit and light him on fire* using an expression in which the imperative operator takes scope over a sentence whose conjuncts can be reversed without affecting that sentence's interpretation. For a logic that has a smooth time capturing such distinctions, see [70, 71]. (I would like to say more about this, but doing so would mean taking a stand on how to resolve the paradoxes I am describing here—something I would like to avoid.)

[8] There has been, for various reasons, a healthy amount of doubt, some of which I am sympathetic to, about the possibility of (valid) imperative inferences in the literature; see [35, 39, 87]. For general defenses of imperative inference, see [23, 85] (although, I should note that I do not find [85]'s defense persuasive). I should note that my own positive view will suppose that imperatives have a logic in only a relatively *weak* sense: their semantics should guarantee that certain constellations of attitudes that are constitutively connected to an agent's acceptance of imperatives display inferential relations to other attitudes that are constitutively connected her acceptance of imperatives. I know of few people who deny that imperatives have a logic in *this* sense.

any well-formed imperative sentence: they must be represented *as embedded* in the regimented language in which their "logical forms" are given.[9]

The clearest case is that of indicative conditional imperatives (CIs for short). The only option $\mathcal{L}_{\texttt{imp}}$ gives for representing a CI is a sentence of the form $!(\phi \rightarrow \psi)$. Such a representation is inadequate in view of the clear consistency of the following sequence of imperatives.[10]

(13) a. If it rains, take the umbrella (but leave the sunglasses)
b. If it doesn't rain, take the sunglasses (but leave the umbrella)
c. Bring both (...since we don't know)

Using only sentences of $\mathcal{L}_{\texttt{imp}}$, we are forced to represent these imperatives using sentences of the form $!(\phi \rightarrow (\psi \wedge \neg\chi))$, $!(\neg\phi \rightarrow (\neg\psi \wedge \chi))$, and $!(\psi \wedge \chi)$. There is, logically, no way to fulfill all of *these* imperatives, so anyone who issues them should be deemed irrational. But a person who issues the imperatives in (13a)–(13c) is not irrational; their advice, in fact, is clearly *sound*. And so $\mathcal{L}_{\texttt{imp}}$ will not do for representing CIs.

If the reader finds this argument wanting, there is no shortage of familiar arguments against wide-scope accounts of conditional imperatives (indeed, the deontic logic literature on so called Contrary-To-Duty Imperatives is thick with them). Here is just one. Consider the following imperatives:

(14) Don't rob Jones
(15) If you rob Jones, do it gently

Using only sentences of $\mathcal{L}_{\texttt{imp}}$, we are forced to represent these imperatives using sentences of the form $!\neg\phi$ and $!(\phi \rightarrow \psi)$. The practical content of the second imperative—how it tells someone to plan—is, thus, strictly *weaker* than the practical

[9] Strictly speaking, the data I appeal to here do not support the claim of irreducible embeddability, since (i) the inability to represent a superficially embedded natural language imperative with a sentence of $\mathcal{L}^{+}_{\texttt{imp}}$ does not imply that (ii) there is no more expressive language that successfully represents that imperative, but represents it as embedded. Indeed, the discussion of Section 5.5 implicitly relies in the falsity of (ii). I hope the reader will forgive some imprecision here in the service of exposition. For clearer examples of embedding, see Section 6.2.

[10] For a substantially similar case involving deontic conditionals, see the Miner Paradox of [42]. Given that it will either rain or not, cases like these may feel puzzling in that they can seem, in a certain light, to support challenges to the validity of modus ponens (understood as a general rule sanctioning the inference of an indicative conditional's consequent if the antecedent is true). But, as [28, 91] emphasize, and as I failed to properly appreciate in my [16], it is possible to formulate the notion of consequence that underwrites valid inference in such a way that insisting on the consistency of (13a)–(13c) is compatible with endorsing a rule of modus ponens.

A common strategy for debunking this sort of example appeals to the idea that the antecedents of the CIs here are elliptical for sentences containing reference to the knowledge of the addressee. That is to say, (13a), on the salient reading, is actually given the reading *if it rains and you know it, take the umbrella*; if it is not given this reading, (13a)–(13c) are inconsistent. This is something I just deny. The salient reading of the CIs in question is not one that makes explicit *reference* to anyone's epistemic states. Insofar as the salient reading yields a consistent interpretation of (13a)–(13c), assigning (13a)–(13c) a consistent interpretation does not require positing covert reference to anyone's epistemic states in the antecedents of indicative conditionals. (For more direct arguments against the strategy I am considering here, see [11].)

content of the former.[11] That is hardly intuitive: even when following (14), (15) can hardly be said to give useless or redundant advice about how to plan.

None of this should be surprising, in view of the fact that is a linguistic datum that CIs (in English, anyway) are genuine indicative conditionals (conditionals being thus typed according to their antecedents, rather than their consequents; see [67, 68]). Indeed, we observe more or less identical phenomena with indicative deontic conditionals, a fact that is generally (and correctly) accepted to render a wide-scope treatment of such conditionals unworkable.

Given this—even bracketing the evident inadequacy of the "wide-scope" analysis of CIs—we might find ourselves interested in the availability of the prevailing treatments of the semantics of indicative conditionals to give an account of the meaning of CIs. Other things being equal, a unified treatment of indicative conditionals would be desirable. But all standard semantic treatments of the indicative conditional (here I'm referring to [46, 51, 72, 73]), though they vary considerably, would seem available only on the assumption that the consequent of the indicative conditional *contributes its truth condition* to the computation of the meaning of the entire conditional.

What is the upshot of all of these considerations?

2.4 Maybe, Just Maybe, Imperatives Have Truth Conditions?

The truth-conditional paradigm has clear advantages, the most obvious being the familiarity and ease with which it can aspire to account for the three facts described just above. This is not surprising. As Seth Yalcin has written about truth-conditional approaches to the semantics of epistemic modality,

> [A truth-conditional strategy] would be nice. For it seems like it would mean less work. It would let us take the view that sentences like these...are not special. It would let us apply to these sentences whatever semantic and pragmatic explanatory strategies we already apply to other uncontroversially descriptive, fact-describing discourse [[90], 295].

And so there is something at least prima facie attractive about trying to *extend the truth-conditional paradigm* to account for our three facts about imperatives. This is not to say that such an extension is automatic. It is more like an expression of optimism. If (*if!*) we can assign imperatives the *right* truth conditions, we may be able to make use of truth-conditional understandings of consistency and consequence, rather than devising new such notions for imperatives. And if we are very lucky, we may be able to appropriate an off-the-shelf semantics for CIs.

To fully appreciate the appeal of this course, it helps to see why the alternatives are unpalatable. In the next section, I will quickly overview two natural ways of going

[11] To forestall a possible response: it might be thought that this claim cannot be right, on the grounds that (i) $!(\phi \rightarrow \psi)$ and $!(\neg\phi \vee \psi)$ are "equivalent", (ii) imperatives of the form $!(\neg\phi \vee \psi)$ have *permission* content: disjunctive imperatives permit either disjunct, and this is why the Ross inference is bad. Whether or not this is true for disjunctive imperatives, it is irrelevant here. Clearly (15) does not permit robbing Jones gently.

non-truth-conditional for imperatives. Both ways are beset with difficulties. This owes largely to their departure from the canonical modes of semantic explanation familiar from truth-conditional semantics.

2.5 The Wages of Non-Truth-Conditionality

Unsurprisingly—since the idea that imperatives have truth conditions seems like an instant non-starter—approaches to imperatives in the philosophical and linguistic literature have largely rejected the notion that imperatives might have truth conditions.

2.5.1 Satisfaction-Conditional Approaches?

Older approaches to imperative logic, spurred largely by R. M. Hare's interest in building a Non-Cognitivist semantics for normative language on the back of a semantics for imperatives, take it for granted that imperatives lack truth conditions, instead aiming to explicate the semantics of imperatives with altogether *different semantic primitives* (e.g., compliance and non-compliance, satisfaction and non-satisfaction, fulfillment and non-fulfillment, bindingness and non-bindingness) (see [12, 24, 32–34, 41, 48, 70]).

Let us leave these authors behind and consider the most prominent contemporary view in this tradition. Peter Vranas has recently taken up this sort of approach in a series of recent papers (see especially his [84–86]). The critical difference in Vranas' approach is that his semantics for imperatives is *three-valued*, rather than two-valued; an imperative, in Vranas' framework, can be *satisfied*, *violated*, or *avoided*. This is an improvement over two-valued approaches (on which the designated semantic values are, e.g., satisfied/unsatisfied), since it allows a treatment of CIs that does not collapse into wide-scope-ism.[12] A CI like the following is satisfied if you turn the A/C on and you shut the window, violated if you turn the A/C on and you do not shut the window, and avoided if you do not turn the A/C on.

(16) If you turn the A/C on, shut the window!

This is nice, so far as it goes. There are, however, downsides.

First, a rather cheap, but nevertheless apposite, observation. As the complexity of Vranas' work attests, pursuing this project in a serious way requires an extraordinary expenditure of effort and ingenuity. For someone interested in the semantics of imperatives in natural language, it is, in the absence of compelling reasons to favor Vranas' approach, fair to count this as a strike against it. Supposing we come up with a simpler approach that is continuous with standard truth-conditional semantics and that

[12] Two-valued approaches collapse into wide-scope-ism because a CI like (16) is either satisfied or not. Obviously it is not satisfied only if: you turn the A/C on and do not shut the window. So, since the semantics is bivalent, it is satisfied only if: you either do not turn the A/C on or do shut the window. In other words, it has the same satisfaction-conditions as an imperative of the form !(*you turn the A/C on* → *you shut the window*).

does as well as Vranas' account with respect to a class of core phenomena, we would have a good, if defeasible, reason to prefer the simpler, more familiar approach. As a corollary to this, supposing we presently lack a simpler approach, but find the task of learning to theorize about and within Vranas' semantics daunting, we have a good, if defeasible, reason to go looking for a simpler approach.

Second, with new semantic primitives come, at a minimum, new notions of consistency and inconsistency. On Vranas' semantics, two imperatives are consistent if it is possible to not violate both (if, for some possible circumstance w, neither imperative is violated in w) [[84], 546]. This notion of consistency is objectionable for two reasons. First, since permission sentences (sentences of the form ¡ϕ) *cannot be violated* (indeed, lack anything like satisfaction or compliance conditions), we gain no purchase on the *inconsistency* of imperatives and contrary permission sentences (e.g., !ϕ and ¡$\neg\phi$).[13] Second, Vranas' notion of consistency leads to bizarre predictions about CIs. For instance, contrary to apparent fact, (16) and (17) are predicted consistent, on Vranas' semantics.

(17) If you turn the A/C on, do not shut the window!

According to Vranas, (16) is violated iff you turn the A/C on and do not shut the window. (17) is violated iff you turn the A/C on and shut the window. But it is possible to avoid violating both: simply do not turn the A/C on! And so, for Vranas, these CIs are consistent.

Obviously, this is a problematic prediction, at least for natural language CIs: (16) and (17) are inconsistent in all of the ways that count (cf. Section 2.1). Compare the following sequence of deontic conditionals, which sounds, for all the world, like a contradiction.

(18) If you turn the A/C on, you must shut the window!
(19) If you turn the A/C on, you must not shut the window!

More generally, any two indicative conditionals with logically inconsistent consequents are plainly inconsistent (which is one reason that philosophers and semanticists have rightly soured on the material conditional analysis of indicative conditionals).

Suppose this worry could be finessed.[14] Even so, a worry would remain. In giving up the familiar, truth-conditional notion of consistency for imperatives, we may escape the frying pan, but nevertheless set ourselves up for an eventual fire. There is a cost to this level of boldness: it requires trailblazing. Trailblazing—especially

[13] This criticism also arguably applies to the Satisfaction-Conditional theory of [23] (although certain complexities Fox's theory, in particular his treatment of Free Choice Permission, make the argument more complex). I am grateful to an anonymous referee for bringing this reference to my attention.

[14] See [[84], 547] for an argument that the sentences should be regarded as consistent (best as I can tell, the argument is a restatement of the claim that the sentences should be regarded as consistent because it is possible not to violate both). In reply, I am inclined to simply insist on the intuition—there is no argument that could persuade me of the consistency of (16) and (17). They are inconsistent in all of the ways that count (cf. Section 2.1).

in terrain as unfamiliar as that defined by imperatives—is difficult and liable to error (and the biting of such unpalatable bullets as this). For this reason, other things being equal, there ought to be a heavy methodological presumption in favor of work that is conservative with respect to major paradigms of semantic theorizing (like, e.g., the truth-conditional paradigm). Considered by itself, this is, admittedly, an extraordinarily weak reason to *dislike* Vranas' theory (indeed, those more inclined than me to appreciate novelty in semantic theorizing may be inclined to regard it as having neutral, or even positive, valence). However, I think it is a fairly strong reason in favor of (i) the *search* for something a little less novel, a little more familiar, (ii) a preference for *working with* simpler theories that are continuous with respect to the theoretically familiar and well-trodden (at least when they can be found).

Finally, Vranas' analysis is of dubious value for the study of imperatives in natural language. Even if this is not his target—and it is most definitely not—it is fair to ask about the theoretical value of developing a detailed semantics for CIs in a highly regimented formal language that has no clear application to theorizing about the semantics of CIs in natural language. CIs in natural language are, after all, *indicative* conditionals (albeit with imperative consequents). As the semantics of the material conditional reflects little of what makes indicative conditionals with declarative consequents so semantically interesting, similarly, a three-valued semantics for a quasi-material conditional (such that it is neither true nor false when its antecedent is false) will, it is likely, reflect little of what makes indicative conditionals with imperative consequents so semantically interesting.[15]

2.5.2 Expressivism?

There are, of course, other options for the fan of imperative non-truth-conditionality. Most obviously, there are so called "Expressivist" or "Speech Act" forms of Non-Cognitivism (classic references are [8, 9, 26, 27]). Such approaches stand in contrast to the Vranas-style approach, which assigns imperatives semantic values of familiar type (a satisfaction condition, which is simply a material conditional whose truth-value is undefined when its antecedent is false—i.e., a *proposition*) but *restyles the relation* between an imperative sentence and its semantic value (not as giving the imperative's truth condition, rather as giving the imperative's satisfaction or compliance condition). Expressivist and Speech Act theories, instead, assign imperatives semantic values of *non-familiar type*—a non-cognitive state of mind (e.g., desire) concerning the actions of the addressee, or a non-representational speech act (e.g., a

[15] Possibly the study of Vranas' imperative conditional could be motivated by appeal to the claim that the logic of this conditional is a topic of sui generis interest. I am doubtful. Imperatives do have a logic (and, like Vranas [84], I think there are valid imperative inferences). So far as I can tell, however, there are no logical or inferential phenomena about imperatives that are both interesting and proprietary to them. The phenomena Vranas cites (like, for instance, the fact that conditional imperatives can be avoided) tend only to loom large in the context of an approach, like his, that takes satisfaction conditions to be an appealing way to theorize about the semantics and logic of imperatives. Since I do not think this, I do not find the phenomena in question particularly interesting.

command)—while keeping the relation between imperatives and their semantic values the same as the familiar relation between declaratives and their semantic values, i.e., as *expression* or *denotation*.

Sundry forms of Expressivist or Speech Act Non-Cognitivism are, of course, possible (and many, indeed, are actual). For our purposes, a single illustration will do. The most well-developed semantics in this tradition, Gibbard's, assigns atomic sentences—declarative normative sentences, in particular, although the idea has a clear extension to imperatives—sets of psychological entities (plans, in the case of declarative normative sentences) as their semantic values, and leaves the rest of truth-conditional semantics untouched. Formally, the connectives still receive their normal Boolean interpretations ($\neg$ expresses $'$, $\wedge$ expresses $\cap$, and so on), and logical consequence is still interpreted as the inclusion of one content within another (so that ϕ entails ψ iff the semantic value of ϕ is a subset of that of ψ). Theoretically, the relationship between sentences and semantic values is as before: sentences express their semantic values. Different forms of Expressivist or Speech Act Non-Cognitivism can be had by varying the kinds of objects to which sentences are mapped in interpretation (although generally these are objects that will determine, in more or less direct fashion, the way those sentences are *used*); this may subsequently warrant revision of the apparatus for compositionally computing the semantic values of complex pieces of syntax, as well as the attendant logical notions.

Despite their differences, there are well-known problems afflicting all such theories. Although the problems are heterogenous in nature, they are standardly grouped under the heading of the "Frege-Geach Problem" (for classic discussions, see [25, 69]). I will mention one. Whatever your favored ontology of non-truth-conditional meanings—speech acts, states of mind, etc.—and whatever elements of that ontology your theory assigns to, e.g., (3a) and (3b), it will be hard for your theory to replicate the explanation of their inconsistency given by the truth-conditional account. For the truth-conditional account, if things go well, will explain their consistency by appealing to a contradiction in the metalanguage; she will assign to (3a) and (3b) truth conditions such that supposing those conditions compatible *generates a contradiction*. As [65, 66] discusses in detail, making use of the sorts of objects that populate non-truth-conditional theories of meaning *prima facie* requires relinquishing this canonical way of explaining object-language inconsistency.[16] For there is *never* anything *contradictory* about supposing an agent expresses two speech acts α and β (or states of mind α and β), regardless of what α and β are: it is always logically possible for an agent to be exceptionally inconsistent.[17]

[16] Schroeder focuses his objections on Expressivist accounts of normative language, but they extend straightforwardly to both Expressivist and speech act accounts of imperatives.

[17] There are serious costs to relinquishing the canonical way of explaining object-language inconsistency—costs that far outstrip its canonicity. You can, for instance, forget about *reductio ad absurdum* reasoning in the metalanguage about object-language inconsistency. For the same reason, it seems that the familiar Tarskian technique of exploiting disquotation—e.g., inferring from the fact that '$a \wedge b$' holds at X the fact that both 'a' and 'b' hold at X—as a tool of proof in the metalanguage will be unavailable. If X is a state of mind, disquotation in this sense is simply *invalid*.

The task is particularly difficult for sentences of the form $!\phi$ and $¡\neg\phi$. Any theorist who tries to explain the inconsistency of these sentences by appeal to the mental states or speech acts they express runs headlong into the most difficult version of the Frege-Geach problem: explaining the inconsistency of attitudes (e.g., disapproving of and tolerating $\neg\phi$) that are provably non-inter-definable (see [65, 66]). It would, at a minimum, be nice to be able to do some ground-level theorizing about the semantics of imperatives without having to resolve the most intractable problem in contemporary meta-ethics.[18]

That is a bit brisk, I realize, and I do not pretend to have given such theories their due here. My point is only to emphasize the prima facie attractiveness of the truth-conditional approach to imperatives. And my overall project in this paper in no way depends on a comprehensive critique of extant non-truth-conditional approaches to theorizing about the meaning of imperatives. Rather, my project here is to sketch a view that retains all the theoretical advantages truth-conditional approaches, while avoiding all of their prima facie objectionable commitments—a view, in other words, that renders the appeal of non-truth-conditional strategies otiose (and so also manages to avoid all of their prima facie objectionable commitments).

3 Two Kinds of Cognitivism

How could such a thing be possible? I will get there eventually. First, though, I will run through a couple different "Cognitivist" semantics for the imperative. After settling on the superior version, and showing how it can account for all of the facts around which we have built our discussion, I will lay out the reasons for being skeptical about Cognitivist semantics, in general, for imperatives. I will ultimately state a semantics that retains the theoretical advantages of the version of Cognitivism I develop, but sheds its baggage.

For reasons that will become clear, I use the label 'Cognitivism' to refer to the view that imperatives express (or mean) propositions—the view that the semantic value of an imperative is a proposition. Cognitivism about imperatives comes in many varieties, depending on the proposition that is thought to give the meaning of an imperative.

3.1 Explicit Performative Cognitivism

Maybe the most well-known version of Cognitivism about imperatives is found in the work of David Lewis [50, 54] and is recently defended (if half-heartedly) in Josh Parsons [57]. Lewis holds, for complicated reasons that I will skirt here, that the

[18] I have argued elsewhere [15] that the sort of semantics developed in this paper, designed in part to skirt the Frege-Geach problem, has a rightful claim to the Expressivist mantle, despite declining to assign speech acts or states of mind as the meanings of sentences. In light of this, here is a more precise statement of my claim here: Expressivism, as ordinarily understood, has a hard time with the Frege-Geach Problem. The issues here are complex, but see my [15] for discussion.

semantic value of any sentence can be factored into a *propositional content* and something like a *force indicator*; such a factored representation will give that sentence's conventional meaning (and will also explain the speech act for which the sentence is typically used by a speaker in a discourse). For declaratives, the procedure is obvious enough: declaratives are associated with a speaker's use of assertoric force, the object of which is the proposition that is ordinarily thought to be that sentence's semantic value.

(20) The cat is on the mat
$\rightarrow \texttt{assert}(x)$(that the cat is on the mat)

The value of x is fixed by the context of utterance; relative to a context c, x gives the speaker of c. The semantic value of a sentence like 'the cat is on the mat' is given by the proposition *the speaker asserts that the cat is on the mat*. Such an analysis for declaratives would be aptly termed an Explicit Performative analysis, since the meaning of a declarative ϕ is identified with the proposition expressed by a performative sentence announcing its performance of the speech act that ϕ simply expresses.

Imperatives, for Lewis, work in much the same way.

(21) Shut the window!
$\rightarrow \texttt{command}(x)(y$ shuts the window)

Here the value of x and y are fixed by the context of utterance; relative to a context c, x gives c's speaker, while y gives c's addressee. The semantic value of an imperative like 'shut the window!' is thus given by the proposition *the speaker commands that the addressee shuts the window*. Such an analysis would be aptly termed Explicit Performative Cognitivism (hereafter EPC) for imperatives.

The Good According to [57], EPC does well with respect to at least Fact 2. Since it assigns imperative sentences truth conditions, it simply inherits a truth-conditional semantics' account of those sentences' logical profile. That (22a) is felt to entail (22b) is explained by the fact that (22a)'s performative "translation" (23a) truth-conditionally entails (22b)'s performative "translation" (23b).

(22) a. Attack at dawn and take no prisoners!
b. Attack at dawn!
(23) a. x commands that y attacks at dawn and takes no prisoners!
b. x commands that y attacks at dawn!

In what sense does (23a) truth-conditionally entail (23b)? It depends on how (23a) is understood: does the command verb scope over a conjunction—$\approx \texttt{command}(x)(p \wedge q)$—or is (23a) itself a conjunction of two command-sentences—$\approx \texttt{command}(x)(p) \wedge \texttt{command}(x)(q)$? If the former, then the relevant entailments will not generally be truth-conditional in nature. As I alluded to above, *command* is, at least to some degree, a non-extensional verb. If an individual commands $(p \wedge q)$, she does not, as a matter of pure logic, thereby command p; although there is a sense in which she is *committed* to commanding p, it is at least logically possible that one can command $(p \wedge q)$, fail to see that the content of her command implies p, and thereby fail to command p. Better, I think, for EPC to understand the

imperative (22a) as having a form something like $(!\phi \wedge !\psi)$ and render its performative translation as a conjunction of command-sentences, so that (23b) is a genuine truth-conditional consequence of (23a) (and, hence, (22b) a genuine truth-conditional consequence of (22a)). This requires a small modification of the imperative language (to allow conjunction to scope over imperatives) but presumably such a modification can be tolerated.[19] So far, so good with respect to Fact 2.

EPC's explanation of Fact 3 is more or less automatic: imperatives contribute the proposition they express—the proposition given by their explicitly performative translation—to the computation of the meaning of complex syntactic constructions in which they occur. So the meaning of (24) is given by the proposition that (25).

(24) If it rains, shut the window!
(25) If it rains, *x* commands that *y* shuts the window

This means that EPC, as Parsons also notes, has a smooth time accounting for the felt validity of modus ponens-like imperative inference like the following.

(26) a. If the weather is good, attack at dawn!
b. The weather is good.
c. So, attack at dawn!

EPC's explanation of this inference's felt validity is immediate: given her understanding of CIs, it *is* simply an instance of modus ponens.

(27) a. If the weather is good, *x* commands that *y* attacks at dawn
b. The weather is good.
c. So, *x* commands that *y* attacks at dawn

The Bad Unfortunately, Explicit Performative accounts are generally unworkable, in ways that doom EPC as a possible account of the meaning of imperatives.

Parsons himself lays out two problems that I take to be decisive. Here they are, briefly. First, there is a problem of "unwanted validities". EPC straightforwardly predicts the following inference valid.

(28) a. Attack at dawn!
b. Someone commands something.

Such an inference is, Parsons suggests, invalid for the same reason that the following inference (predicted valid by the Explicit Performative analysis for declaratives!) is invalid:

(29) a. It is sunny.
b. Someone asserts something.

[19] See my [13] for an argument that this sort of wide-scope conjunction is obligatory.

Second—and, to my mind, more worrying—is a problem of "unwanted consistencies". EPC predicts the following (obviously inconsistent) sentences consistent.

(30) a. Attack at dawn!
b. Don't attack at dawn!

But their performative translations are obviously consistent:

(31) a. I command that you attack at dawn!
b. I command that you do not attack at dawn!

Parsons claims that "it is not inconsistent to report that someone (even oneself) has commanded inconsistently" (53). This claim is a bit awkward: of course it *is* inconsistent for someone to jointly assert (31a) and (31b). One takes Parsons' point to be that the performative sentences *themselves* are logically consistent (jointly satisfiable), while their corresponding imperatives are not. This in fact deprives EPC of one of the signature advantages of truth-conditional theories of meaning for the imperative: it cannot give a canonical explanation of inconsistency (where a canonical explanation of inconsistency is one that is associated with a contradiction in the metalanguage). So EPC has trouble—as much trouble, in fact, as any of its non-truth-conditional rivals—with Fact 1.

3.2 Modal Cognitivism

Fortunately for the fan of the truth-conditional paradigm, there is a better option out there: the *Modal* Analysis. The Modal Analysis is, to my mind, by far the most plausible form of Cognitivism for imperatives (and it is certainly the form of Cognitivism with the most empirical support in the linguistics literature; see [2, 31, 40, 67, 68]). This is not altogether surprising, given the obviously tight connection between imperatives and deontic modals (e.g., *must*, *should*, *ought*, etc.).[20]

In a nutshell, on the Modal analysis, there is an identification in meaning between (32) and (33)

(32) Attack at dawn!
(33) You should attack at dawn!

[20]For this reason, the Modal Analysis of imperatives has frequently been seen as a contender in the historical discussion of imperatives. [29] in fact dates the Modal Analysis to Kant's *Groundwork*! The Modal Analysis was certainly seen as a contender in the Hare-era debate on imperatives (see [24, 33]), but was generally dismissed on the grounds that imperatives, unlike deontic modals, were "two-valued" with regard to their prejacents. As Geach explained, "In answer to a request for orders, 'am I to do P?', only two answers are possible—'do P' and 'do not do P', which are contradictories. No other order is a direct answer to the question; and to say 'you may either do P or not' would not be an answer but a refusal to answer—I was asked for an order and I refuse to give any. But there are three relevant answers to the moral enquiry 'ought I to do P?'" [[24], 49]. Geach's observation is questionable—I think 'it is up to you' counts as a perfectly good answer to the question 'what am I to do?' (unless the interrogator is supposing that only an order will do). Even if it were right, it would show only that some *imperative questions*—some questions of the form *am I to do P?*—were distinct in meaning from *deontic questions*—questions of the form *ought I to do P?* This is irrelevant to the question of whether the answers to imperative questions are distinct in meaning from the answers to deontic questions.

In what follows, I will generally say that Modal Cognitivism represents the *logical form* of an imperative with the corresponding modal sentence. (The specific modal I have chosen, *should*, is irrelevant to my argument; everything I say could be recast using *ought* or *must*.)

Somewhat ironically (since the most famous version of the Explicit Performative analysis is due to David Lewis [50, 55], the most famous version of the Modal Analysis is also due to David Lewis [53].[21] Letting '$\Box$' symbolize an operator expressing deontic necessity, and $[\![\cdot]\!]$ a function mapping a sentence to its semantic value, Lewis' analysis has the following parts:

- $[\![!\phi]\!] = [\![\Box\phi]\!]$
- $!\phi$ is true relative to a *sphere of permissibility* (a set of situations that can be permissibly realized) S (notation: $[\![!\phi]\!]^S = 1$) iff $[\![\Box\phi]\!]^S = 1$ iff ϕ holds throughout S (notation: $S \subseteq \|\phi\|$), where $||\phi||$ is the set of possible worlds at which classically holds.
- The rules that govern the use of an imperative are such that, if it is false at the sphere of permissibility at the time of utterance, "the sphere adjusts itself, if possible, to make the [speaker's] sentence true" (164). (For Lewis, this happens by eliminating certain elements of the permissibility sphere, namely those incompatible with imperative's content. So, the result of updating S with $!\phi$ is $S \cap \|\phi\|$.)

On this version of Lewis' account, the effect of uttering an imperative is presumably the same as the effect of the corresponding explicit performative; both make their complement propositions obligatory. Nevertheless, its logical form is distinct: the logical form of an imperative, unlike the logical form of the corresponding explicit performative, is modal in nature.

I state Lewis' account of the semantics of imperatives and their corresponding deontic necessities only to give a concrete idea of how such an account is actually supposed to work. As is well known, any account of deontic necessity modals that treats them as universal quantifiers—as Lewis' simple account in fact does—is subject to the sorts of paradoxes I described above (on this point, see especially [10]). Hence, any account of imperatives which assimilates them to universal quantifiers will be subject to the same paradoxes. Regardless of one's attitude toward these paradoxes, the semantics of deontic necessity modals is much more complicated than this sketch would suggest; at a minimum, such modals are interpreted with respect to some sort of goodness ranking or preference ordering, a detail that Lewis' account suppresses (cf. [43–45]).

None of this, however, is strictly relevant to the prospects for Modal Cognitivism about imperatives. Whatever the correct semantics for deontic necessity modals, and whatever resolution of the semantic paradoxes for deontic logic we opt for, can be appropriated by Modal Cognitivism. Modal Cognitivism itself is not hostage to any specific account of the semantics of deontic necessity modals. It says only that the

[21] See also [2, 4, 31, 67].

correct account of their semantics can do double duty as an account of the semantics of imperatives.

The Good Modal Cognitivism shares certain virtues of the Explicit Performative Cognitivism (and adds more besides). Since an imperative's logical profile is the same as that of its modal "translation", Modal Cognitivism offers an automatic (and rather appealing) account of Fact 2. That (22a) entails (22b) can be explained by the fact that (22a)'s modal "translation" (34a) truth-conditionally entails (22b)'s modal "translation" (34b).

(34) a. You should attack at dawn and take no prisoners!
b. You should attack at dawn!

No need to modify the basic imperative language (to allow conjunction to scope over imperative sentences) to account for this fact: on any reasonable treatment of the semantics of deontic logic, (34b) will be a *logical consequence* of (34b).

Similarly, Modal Cognitivism, since it is a form of Cognitivism, offers an automatic explanation of Fact 3: imperatives contribute the proposition they express—the proposition given by their modal "translation"—to the computation of the meaning of complex syntactic constructions in which they occur. So the meaning of (35) is identified with the meaning of (36).[22]

(35) If it rains, shut the window!
(36) If it rains, you should shut the window!

Again, Modal Cognitivism is not hostage to any specific account of the semantics of indicative imperative or deontic conditionals. Whatever the correct account of such conditionals—as Variably Strict conditionals (à la [51, 72]), or as generalized quantifiers where the *if*-clause functions solely to restrict the domain of the quantifier (à la [46])—Modal Cognitivism can point to it and say: *there, that's what I want to say about the semantics of indicative imperative conditionals*.

Modal Cognitivism, like EPC, explains the validity of the modus ponens-like imperative inference in (27) by appeal to the truth-conditional validity of (37).

(37) a. If the weather is good, you should attack at dawn
b. The weather is good
c. So, you you should attack at dawn

Do indicative deontic conditionals support modus ponens? Yes![23] So CIs will support it too.

The Better Modal Cognitivism obviously succeeds where EPC failed. There is no question of inference (28) being licensed, since inference (38) is truth-conditionally invalid.

[22] Exactly this fact is enthusiastically exploited by [67, 68] to support a modal account of imperative semantics.

[23] Something close enough, anyway. See [42, 91] for discussion.

(38) a. You should attack at dawn
b. Someone commands something

The problem of unwanted validities, at least as it afflicted EPC, disappears on Modal Cognitivism.

Ditto the problem of unwanted consistencies. Since (39a) and (39b) are truth-conditionally inconsistent, (30a) and (30b) are correctly predicted inconsistent.[24]

(39) a. You should attack at dawn!
b. You should not attack at dawn!

As a corollary, Modal Cognitivism directly explains the intuitive inconsistency of imperatives and their contrary permission-grants. The logical form of an imperative of the form $!\phi$ is given with a modal formula of the form $\Box\phi$. If we extend Modal Cognitivism to permission-grants, we can assign a permission-grant of the form $¡\neg\phi$ a logical form $\Diamond\neg\phi$. The modal "translation" of $!\phi$, $\Box\phi$, is straightforwardly inconsistent with the modal "translation" of $¡\neg\phi$, $\Diamond\neg\phi$. While EPC ran headlong into the hard version of the Frege-Geach problem, Modal Cognitivism skirts it effortlessly.

Moreover, Modal Cognitivism is attached to an account on which the relationship between imperatives and their corresponding modals is explicitly laid out: an imperative $!\phi$ typically functions to make its corresponding deontic necessity claim $\Box\phi$ true.[25] This is precisely the sense in which imperatives are typically understood to *command*.

In sum, Modal Cognitivism has *all of the virtues* and *none of the vices* of EPC. These virtues are, we've seen, considerable, at least given our goal of account for Facts 1, 2, and 3. Maybe, just maybe, imperatives have truth conditions after all?

4 Against Cognitivism

Of course, Modal Cognitivism comes with a serious drawback: it identifies the semantic content of an imperative with that of a modal sentence: both express, as a matter of their semantics, *propositions*. Why, exactly, is this a drawback? And how serious is it? Very serious, indeed, as I will explain in this section.

4.1 Truth

Parsons [57] worries about Cognitivism on the grounds that imperatives, at first blush, "are not in the business of describing the world, but of telling people what to do. They cannot therefore be true or false—for to be true is to succeed in the business of describing the world (and to be false is to fail)" (49). Cutting against this, however, is a kind of flexible *explanatory pragmatism* about truth. Suppose assigning imperatives truth conditions is of explanatory value (in, e.g., accounting for their logical profile)—value that cannot be replicated by any non-truth-conditional account.

[24] I assume, as is plausible, that (30b) can be represented with an imperative of the form $!\neg\phi$.

[25] On the importance of accounting for this dimension of imperative meaning, see [59].

In light of this fact, running afoul of this specific philosophers' platitude—that things that are true or false can be used to describe the world—might seem a small price to pay.

Well, even a healthy pragmatism has its limits. At the very least, the sort of pragmatism about truth that is being encouraged should allow us to say some true things—things we would not otherwise be able to say—about specific imperatives. But does it? That is unclear. It is not just manners or style that lead us to blanch when we are asked to ascribe truth or falsity to an imperative. Such ascriptions are *ungrammatical*. How is an explanation of the inconsistency of, say, $!\phi$ and $¡\neg\phi$ supposed to get going, if we cannot *grammatically express* the reductio assumption in our metalanguage?

(40) ??Suppose it is true that go to your room and that you may stay up. Then...

What good are truth conditions for imperatives if they do not enable us to explain, in terms that are, at a minimum, *intelligible* to us, the things we are interested in explaining?

Cognitivism, in other words, claims its gains a bit too quickly. It is not enough to simply have a *formally adequate* characterization of (in)consistency and consequence for an imperative language—one that correctly sorts sets of sentences into consistent and inconsistent piles, and arguments into valid and invalid piles. For all we have seen, there is no disputing that Modal Cognitivism, at least, gets us that: two imperatives are consistent just if their modal "translations" are, while an imperative argument is valid just if its modal "translation" is. But such characterizations, in addition to being formally adequate, need to meet certain demands of *theoretical adequacy*. In particular, they need to *explain* that which they claim to explain (particularly when their sole claim against the incredulous stare—*how could imperatives possibly have truth conditions?!*—is their adverted explanatory value).

In short, Modal Cognitivism gives us the tools for a formally adequate characterization of imperative (in)consistency and consequence. But it does not tell us how to marry these characterizations to a *theory* of imperative (in)consistency and consequence. Modal Cognitivism, for all its advantages, lacks such a theory. Given the very basic fact that we cannot even *entertain the thought that* an imperative is true, it is hard to see how such a theory might be intelligibly stated. And this fundamentally calls into question Modal Cognitivism's claimed explanatory value.

4.2 Function

On top of this, there are, I will now argue, strong reasons to think that imperatives are distinct in semantic value from their corresponding modal sentences. My argument trades on the following sort of inference (which I motivate shortly).

1. Imperative language has a non-representational function
2. Imperative language has a non-propositional semantics

A tacit commitment to this sort of connection between semantic types (propositions) and cognitive functions (representation) is, I conjecture, the central reason most of us are loath to think imperatives might have propositions as their semantic

values. We do not want to incur the significant burden of explaining why, if imperatives express, as a matter of their semantics, propositions:

- It does not make sense to say that an imperative is (or says something) true or false
- One cannot, in normal contexts, perform an assertion with an imperative
- One cannot ascribe belief in (or assertion of) an imperative

Such facts are sufficiently well-attested by the following examples:

(41) A: Go to your room!
#B: That's true! [#That's false!]
(42) A: Shut the window!
#B: You're right, I didn't realize that!
(43) #Sam asserted that go to your room
(44) #Sam thought that shut the window

The truism underlying our aversion to a propositional semantics for imperatives is, plausibly, just this: *if imperatives had a propositional semantics, we would expect them to have a representational function*.[26]

One need hardly be a disciple of Wittgenstein or Austin to want a semantics for a sentence to account for certain facts about its function—facts like:

- What that sentence is conventionally used, by a speaker, to do
- What acceptance of that sentence canonically consists in

Indeed, such a connection is very much part of the way both philosophers and linguists think about the relationship between (i) a sentence's clause-type (often referred to as its "mood"[27]), (ii) its semantic type, and (iii) its characteristic use and cognitive profile. I will use the example of declaratives to illustrate.

As is attested by a great deal of linguistic work on clause-type, clauses of the declarative type are standardly used to make assertions, resolve issues, express beliefs, and so on. The fact that clauses of the declarative type are standardly used to make assertions, express beliefs, and resolve issues is elegantly explained by the fact that their meanings encode a certain picture of the world (for a good overview, see [58]). Propositions, at a minimum, *encode a picture of a way the world might be*. When someone utters a sentence whose meaning is a proposition, they *proffer*—literally, *offer for acceptance*—the proposition to their audience. In proffering a proposition to an audience—offering that proposition for acceptance—a speaker

[26] Strictly speaking, the tests suggested here for whether a sentence has a representational function are not decisive. There are *many* philosophers who think that, although things like indicative conditionals, epistemic modals, and moral claims certainly can be called true/false and can be asserted and believed, they nevertheless lack both a representational use and a propositional semantics. Ultimately, I think the proper way to draw the representational/non-representational distinction for a sentence appeals to cognitive phenomena—whether, say, Bayesian probability functions are defined over its semantic value—rather than crude linguistic tests of the sort suggested here.

[27] Not to be confused with the syntactic mood of its main verb. The verbal mood of an interrogative is typically the same as its (declarative) answers. But interrogatives and declaratives are different kinds of clauses.

is naturally taken by her audience to endorse that proposition's representation of the world. All of this is familiar, more or less, from Stalnaker's classic account of assertion, according to which "The purpose of expressing propositions is to [distinguish among alternative possible ways that things may be]" [74].

It is easy to see that this creates trouble for a view on which imperatives mean propositions. Whatever imperatives are standardly used to do—commanding or, perhaps, something else—it is definitely *not that*. First, a platitude about communication: in uttering a sentence ϕ with meaning M, a speaker typically proffers M to her audience. If, then, imperatives have propositions as their meanings, utterances of imperatives will proffer propositions to an audience. It will thus appear wholly mysterious that imperatives cannot partake in the speech acts of assertion, belief-expression, and issue-resolution.

Probably, then, imperatives do not, as a matter of their semantics, function to proffer propositions to an audience. They proffer something else: perhaps prescriptions (see [84]), perhaps actions (see [5, 58, 59]), perhaps speech acts, and perhaps (as I will argue) something else entirely.

4.3 Signpost

Let's review our progress. Assigning imperatives truth conditions, as I argued in Section 2, seems like a good strategy for handling Facts 1, 2, and 3 about the semantics of imperatives. This point was strengthened by the difficulties that beset the several non-truth-conditional approaches which we surveyed (Section 2.5). It was further strengthened by the apparent formal adequacy of one version of the truth-conditional strategy: Modal Cognitivism. But it was weakened by the difficulties that beset the most obvious way of implementing the truth-conditional-strategy (which involved saying that imperatives express, as a matter of their semantics, propositions, and used these propositions to account for Facts 1, 2, and 3).

Modal Cognitivism certainly has a leg up over Vranas' satisfaction-conditional approach, which does not seem even formally adequate with respect to (i) the inconsistency of an imperative of the form $!\phi$ with a contrary permission grant of the form $¡\neg\phi$, (ii) the inconsistency of contrary conditional imperatives.

But compared to the family of views I classed as Expressivist or Speech Act-Theoretic, its status is less clear. Expressivist theories, I mentioned, had trouble generating canonical explanations of inconsistency. For there is *never* anything *contradictory* about supposing an agent expresses two speech acts α and β (or states of mind α and β), regardless of what α and β are. Modal Cognitivism's supposed advantage over such a theory is in its claimed ability to generate canonical explanations of inconsistency: the claimed advantage is not with respect to formal adequacy, but rather theoretical adequacy. But there is a real question about whether the explanations of inconsistency offered by Modal Cognitivism are theoretically adequate at all. And Modal Cognitivism has the further disadvantage, which Expressivism can very easily avoid, of having difficulty explaining why imperative language has a non-representational function.

In short, it is beginning to look like we are forced to choose between bad and worse options for theorizing about the semantics of imperatives. Are our theoretical prospects really this bleak? In the subsequent section, I will argue that they are not.

5 Modal Noncognitivism

On the view I wish to sketch here, which I will rather unimaginatively term *Modal Noncognitivism*, imperatives and their corresponding modal sentences are intimately related. Nevertheless, a modal sentence does not in any sense give the meaning, or logical form, of an imperative. Imperatives express—have as their meaning, or semantic values—*propositional concepts*. An imperative of the form $!\phi$ expresses a function that, when evaluated relative to a certain kind of object—a plan—determines a modal proposition—the proposition that the plan requires the imperative's prejacent, ϕ.

I begin this section by motivating Modal Noncognitivism. Modal Noncognitivism arises, I suggest, out of a natural understanding of the function of imperative sentences: an understanding on which imperatives tell agents how to plan, rather than what to believe. I explain why Modal Noncognitivism does better than its competitors with respect to our three Facts. Finally, I sketch how Modal Noncognitivism could be compositionally implemented in a semantics for indicative conditional imperatives. I close the section by giving reasons to prefer Modal Noncognitivism to two different, well-developed theories of the meaning of imperatives: Paul Portner's [58, 59] account, on which imperatives proffer their prejacents for addition to the addressee's "To-Do List, and Will Starr's [75] account, on which imperatives update a body of preferences so that it comes to prefer the imperative's prejacent to its negation.

5.1 Design: From Function to Semantics

To explain the canonical use of imperatives (and permission sentences), it would be nice to be able to follow the following broadly Expressivist idea:

> **Semantics Explains Function**
> The semantics of imperatives (and permission sentences) makes them well-suited to a specific kind of non-representational function[28]

Here is what I think is the clearest way of doing this. Treat sentences as functioning to resolve *cognitive issues*: questions about what to believe, what to want, what to pursue, and so on. Sentences resolve such questions by encoding (perhaps—and, I will suggest, in fact—semantically) cognitive instructions.[29]

[28] Portner [58, 59] is a forceful defender of this sort of principle.

[29] This sort of idea is inspired by and closely related to the approach to semantic theorizing advocated by [77, 78], and to a somewhat lesser extent the approach of [89, 90].

Function as Cognitive Instruction
The function of a sentence is to resolve a cognitive issue via providing a cognitive instruction

Nothing controversial is intended by this. I focus on cognitive instructions because they occupy a nice middle ground between speaker-intention-oriented and addressee-update-oriented approaches to theorizing about a sentence's function. Speakers select sentences with an eye to realizing a specific communicative intention; *which* sentence a speaker selects will presumably depend on how that sentence instructs her addressee to update. In normal contexts, an addressee may recognize a speaker's communicative intention by exploiting information about what cognitive instruction her utterance is suited for communicating; updating on the speaker's utterance means updating according to the cognitive instruction the speaker is interpreted as intending to communicate.

How does this idea connect up to semantic theorizing? Imperatives, to an approximation, tell agents *what to plan*, while sentences with a representational use tell agents *what to believe*. In other words, an imperative $!\phi$ expresses a property of plans: the property a plan has when it is in line with how the imperative tells the agent to plan. Which property is that? To an approximation, the property a plan has if it is decided on ϕ. Similarly, $¡\neg\phi$ expresses a property of plans. Which property? To an approximation property a plan has when it is tolerant of $\neg\phi$.

So far so good. But we must be careful. Suppose we try to turn the above reflections into a semantics for imperatives. Suppose we do this by saying:

- The semantic value of an imperative $!\phi$ is the property of being decided on ϕ
- The semantic value of a permission $¡\neg\phi$ is the property of being tolerant of $\neg\phi$

This, however, gets us into trouble. It gets us into trouble because we find it difficult say *what is inconsistent* about a plan that is decided on ϕ and one that is tolerant of $\neg\phi$. This is the "strong version of the Frege-Geach problem" that afflicted Expressivism. It is something to avoid.

Opponents of Expressivism claim that this shows that it is mistaken to think that *psychological kinds* or *properties* (like properties of human plans) can function as semantic values (see [65, 66]). The stuff of human psychology is not endowed with the right sorts of characteristics to offer self-sufficient explanations of what is wrong with being decided on ϕ and tolerant of $\neg\phi$. In this, I am in agreement with the opponents of Expressivism. In particular, as I have already stated, I am uncomfortable with the fact that Expressivism (like Modal Cognitivism) seems to require that we relinquish our canonical semantic methodology for explaining object-language inconsistency: the derivation of *a contradiction* in the semantic metalanguage. There is simply nothing *contradictory* about supposing that the psychological properties *being decided on* ϕ and *being tolerant of* $\neg\phi$ are co-instantiated. Instantiating both is, perhaps, irrational, but nevertheless possible.

Nevertheless, I still think it is possible to follow the broadly Expressivist idea with which we began this section, while retaining our canonical semantic

methodology for explaining object-language inconsistency. Here is how I suggest we do this:

- Retain the core of traditional model-theoretic semantics: treat **abstract** or **formal entities** as providing **verdicts** for sentences. Sentences, as a matter of their semantics, place conditions on these formal entities such that: if the entity e meets the condition specified by a sentence ϕ, then ϕ receives a positive verdict relative to e.
- Declaratives place conditions on sets of possible worlds.
  Sets of possible worlds serve as abstract representations of an agent's information.
- Imperative (and permission) sentences place conditions on sets of action-descriptors.
  Sets of action descriptors serve as abstract representations of an agent's plans.

It is important to note that there is *nothing new here*, at least as far as declaratives are concerned. We can continue to understand atomic declaratives as expressing propositions—for simplicity, sets of possible worlds. A possible world is a set of atomic sentences: the set of atoms that are true at that world. So propositions are typed as sets of sets of atomic sentences. An atomic declarative p holds at a set S of possible worlds iff for every $w \in S$, $p \in w$. This can be turned into a recursive characterization of *holding at* as follows.

Holding At
$S \models p$ iff $\forall w \in S : p \in w$
$S \models \neg\phi$ iff $\forall w \in S : \{w\} \not\models \phi$
$S \models (\phi \wedge \psi)$ iff $S \models \phi$ and $S \models \psi$[30]

It is obvious enough that, when ϕ is declarative, $S \models \phi$ iff $S \subseteq ||\phi||$ (where $||\phi||$ is the set of possible worlds at which ϕ classically holds). All that is new is a *cognitive gloss* on this relation between a set of worlds and a declarative: p holds at S iff at every world compatible with S, p, iff S "believes" or "accepts" p. The cognitive gloss reflects some of what, I have argued, a propositional semantics for declaratives already encodes: declaratives proffer a proposition *by way of* proffering a condition on an agent's beliefs—by way of telling an agent *what to believe*. Indeed, the condition on an (abstract representation of an) information-state that a declarative expresses is *derivable from* its ordinary propositional semantics.

[30]Disjunction is defined in terms of negation and conjunction as follows. $S \models (\phi \vee \psi)$ iff $S \models \neg(\neg\phi \wedge \neg\psi)$. This holds iff $\forall w \in S : \{w\} \not\models \neg\phi \wedge \neg\psi$ iff $\forall w \in S : \{w\} \not\models \neg\phi$ or $\{w\} \not\models \neg\psi$ iff $\forall w \in S : \{w\} \models \phi$ or $\{w\} \models \psi$. This is the correct notion of disjunction for the sake of our cognitively oriented semantics: for a disjunction $(\phi \vee \psi)$ to hold ("be known") at S, it is required only that, for every world in S, at least one of ϕ or ψ holds at that world. This is crucially (and correctly) distinct from a notion of disjunction on which $(\phi \vee \psi)$ holds at S iff $S \models \phi$ or $S \models \psi$. Knowing $(\phi \vee \psi)$ does not require knowing at least one of ϕ or ψ; one can know a disjunction without knowing which disjunct is true.

5.2 Basic Semantics for Basic Imperatives

The proposal I am suggesting—to have a handy name, let's refer to it as **Model- and Cognitive-Conditional-Theoretic Semantics** (MCCTS)—is, frankly, not all that interesting for proposition-expressing declaratives. Not so for imperatives.

Here is the nut. The condition C which a sentence places on a semantic parameter π corresponds to a proposition: the proposition that π meets C. Two sentences sentence can thus place *logically incompatible conditions* (as opposed to just rationally incompatible conditions) on a single parameter. So supposing that some parameter π meets two conditions C and C' can, in some cases, generate a genuine contradiction in the metalanguage.

Here is how this works for declaratives. Consider an atomic sentence p and its negation $\neg p$. Suppose that for some sets of worlds S, $S \models p$ and $S \models \neg p$. Then for every world w compatible with S, $p \in w$ and $\{w\} \not\models p$. But then for every world w compatible with S, $p \in w$ and $p \notin w$. That is a **contradiction**. What this contradiction reflects is the fact that an atomic declarative and its negation place logically incompatible conditions on a semantic parameter.

It would be natural to suspect that this sort of technique can be applied only if we are working with "familiar" semantic parameters: parameters that are set-theoretic constructions out of possible worlds. Actually, however, this is not so. The condition C which a sentence places on a semantic parameter π corresponds to the proposition that π meets C. It does not matter what sort of stuff π is made of. All that matters is that we can extract propositions with the right characteristics from the conditions that the relevant sentences place on the relevant parameters.

This is all fairly abstract, so let me make it more concrete by explaining how this technique might work for imperatives. Recall the general idea about imperative semantics proposed in Section 5.1:

- Imperatives tell an agent *what to plan*
- An imperative of the form $!\phi$ does this by encoding, as a matter of its semantics, the property a plan has if it is decided on ϕ. The semantic function of an imperative is to partition plans along such lines.

Recall also our commitment to retaining the core of traditional model-theoretic semantics, by treating abstract entities—for imperatives, sets of action-descriptors, for declaratives sets of possible worlds—as providing verdicts for sentences. Here is how to bring these two ideas together.

Decision as Requirement, Toleration as its Partial Dual Let us say a plan (here and subsequently, to be read in the abstract sense, i.e., as a set of action descriptors) Λ is decided on ϕ, in the relevant sense, iff ϕ is required relative to Λ; the condition an imperative places on a plan for a positive verdict is that the plan make it the case that it requires ϕ. In sum, an imperative $!\phi$ expresses the property a plan Λ has just if Λ requires ϕ:

- $[\![!\phi]\!] = \lambda\Lambda\,.\,\Lambda$ requires ϕ

The meaning of an imperative of the form $!\phi$ is to be identified with just such a function (or, if you like, just such a property). On the other hand, a permission $¡\phi$ expresses the property a plan π has just if Λ allows ϕ.

- $[\![¡\phi]\!] = \lambda\Lambda\,.\,\Lambda$ allows ϕ

The relevant notions are intended as partial duals. What I mean by this is that Λ requires ϕ only if Λ does not allow $\neg\phi$, while Λ allows ϕ only if Λ does not require $\neg\phi$.[31]

Requirement as Modal How can we get a requirement-condition from a set of action descriptors? There are many ways. Here we follow one influential strategy in natural language semantics, due to Kratzer [44], which uses a plan to determine a desirability or preference ordering on possible worlds, and subsequently treats notions like *requirement* and *allowance* as quantificational (or modal) in nature. This is, obviously, a very simplified treatment of preference, but it should be concrete enough for the very simple purposes to which I will be putting it here.

Preferences from Plans We are treating plans as sets of action descriptors. A plan for an agent a, Λ_a thus determines a set of propositions P_{Λ_a}: the set of propositions p such that for some $\alpha \in \Lambda_a$, p is the proposition that a performs α. Such a set of propositions can straightforwardly determine a preorder on possibilities as follows.

$$w \preceq_{\Lambda_a} v \text{ iff } \{p \in P_{\Lambda_a} : v \in p\} \subseteq \{p \in P_{\Lambda_a} : w \in p\}$$

In English, w is at least as good as v, with respect to Λ_a, iff w involves a's performance of at least those actions in Λ_a that a performs in v, and perhaps more besides.

We will assume (for simplicity) that the alternatives that are ranked by such a preference ordering are generally finite. It will always, then, hold that some worlds are minimal (cannot be strictly improved on) with respect to the preference ordering.

$$\exists w : \neg\exists v : v \preceq_{\Lambda_a} w \text{ and } w \not\preceq_{\Lambda_a} v$$

Relative to a plan Λ_a, we will refer to the set of such worlds that are $\preceq_{\Lambda_a}$-minimal as the Λ_a-*best* worlds, and notate this set as min_{Λ_a}. Sometimes the modal's domain is restricted to a set R, in which case we will notate the Λ_a-best R-worlds as $min_{\Lambda_a}(R)$.

[31] There is a familiar worry about treating these properties as full duals, i.e., about turning these 'only if' claims into 'iff' claims. For one way for a plan to fail to allow $\neg\phi$ (without thereby requiring ϕ), as well as to fail to require $\neg\phi$ (without thereby allowing ϕ), is for it to be *undecided* about ϕ (see [21, 65, 66, 79, 80]).

In my view, this is basically parallel to the problem of distinguishing *strong permission that* $\neg\phi$ (allowing $\neg\phi$) from *weak permission that* $\neg\phi$ (being undecided about $\neg\phi$, i.e., failing to require that ϕ, while also failing to "actively" permit that $\neg\phi$). It is an old problem, and one on whose resolution I would like to avoid taking a stand on here. However, a natural idea here would be to imbue plans with more structure, so that the distinction between strong and weak permission can be represented (cf. Yalcin's [90] addition of the sort of structure to information states allowing the representation of the attitude of treating something as an issue).

Quantificational Semantics for Requirement and Allowance As we saw in Section 2.2, the semantics of the deontic notions of requirement and allowance is fraught with complexity and paradox. I do not want to take any stance here on how such issues are to be handled in a semantics for such notions. So I will state the barest bones semantics that allows me to explain what I wish to explain so far as imperatives are concerned. I provisionally suggest the following partial analyses of requirement and allowance with respect to a plan.

- Λ_a requires ϕ only if $\forall w \in min_{\Lambda_a} : \{w\} \models \phi$
- Λ_a allows ϕ only if $\exists w \in min_{\Lambda_a} : \{w\} \models \phi$

By declining to state sufficient conditions for requirement and allowance, I mean to avoid Paradoxes, like the Ross Paradox, which seem to show that, e.g., notions like requirement and obligation have content over and above universal quantification over the best possibilities.

In Sum Our semantics for basic imperatives (and permissions) can be stated very briefly.[32]

- $[\![!\phi]\!] = \lambda \Lambda_a \,.\, 1$ only if $\forall w \in min_{\Lambda_a} : \{w\} \models \phi$
- $[\![¡\phi]\!] = \lambda \Lambda_a \,.\, 1$ only if $\exists w \in min_{\Lambda_a} : \{w\} \models \phi$

An imperative of the form $!\phi$ places the following as a condition on a plan: the plan must prefer ϕ worlds to $\neg\phi$ worlds; a plan receives a positive verdict (i.e., 1) only if it meets this condition. A permission of the form $¡\phi$ places the following as a condition on a plan: the plan must not prefer $\neg\phi$ worlds to ϕ worlds; a plan receives a positive verdict (i.e., 1) only if it meets this condition.[33]

This, I will preliminarily submit (I will have some more to say later), gives (part of) the sense in which an imperative (permission) of the form $!\phi$ ($¡\phi$) encodes, as a matter of its semantics, the property a plan has if it is decided on (allows) ϕ. By encoding such a property, an imperative (permission) tells an agent what/how to plan.

[32] Notice that the λ-notation here is to be read non-standardly (since describing the function using 'only if' instead of 'iff' means that we fail to pick out a unique such function). A λ-expression of the form $\lambda z Fz$ will, therefore, refer to **a** (rather than **the**) function mapping z to Fz.

A referee for this journal has helpfully worried that this leaves the resulting semantics quite incomplete—that I have only described a class of functions that an imperative or permission sentences might denote, rather than the particular functions they do denote, and that the cognitive instruction an utterance of an imperative will, on my view, proffer is, thus, quite indeterminate. Quite so, and I agree that this is a shortcoming of the view as stated. But the view as stated is to be understood provisionally: the semantics I would, in the final instance, endorse *does not* assign families of functions as semantic values; a unique function will be assigned, and the cognitive instruction thus proffered in expressing such a function will be fully determinate. Since stating such a semantics would, of necessity, require taking a stand on the above mentioned paradoxes, I have declined to do so here. Instead, I have opted for stating precise necessary conditions on the function expressed by an imperative. I have done this for the sake of showing that the framework, no matter how it is filled out, will be able to give satisfactory explanations of the phenomena I am interested in here.

[33] For a similar view of the preferential content of imperatives, see [75].

5.3 Inconsistency

It follows pretty much *immediately* that an imperative of the form $!\phi$ and a contrary permission of the form $¡\neg\phi$ are inconsistent.

Proof Suppose for some Λ_a, $\Lambda_a \in [\![!\phi]\!]$ and $\Lambda_a \in [\![¡\neg\phi]\!]$. Then $\forall w \in min_{\Lambda_a} : \{w\} \models \phi$ and $\exists w \in min_{\Lambda_a} : \{w\} \models \neg\phi$. Then $\forall w \in min_{\Lambda_a} : \{w\} \models \phi$ and $\exists w \in min_{\Lambda_a} : \{w\} \not\models \phi$. Contradiction. □

So our semantics is at least formally adequate, at least with respect to the inconsistency of imperatives and contrary permissions.

On Theoretical Adequacy Is it theoretically adequate? That is to say, does it enable us to *explain*, in terms that are intelligible to us, the inconsistency of an imperative and contrary permission? It does. The explanation goes like this. Generally speaking, sentences place conditions on formal parameters: conditions that a parameter meets when the sentence holds at the parameter. Imperatives and permissions place conditions on abstract representations of plans, while declaratives place conditions on abstract representations of information states (which, as we saw, reduce to conditions on possible worlds). A plan Λ meets the condition placed by an imperative of the form $!\phi$ only when Λ requires ϕ, while a plan Λ meets the condition placed by a permission of the form $¡\neg\phi$ only when Λ allows $\neg\phi$. Supposing there is any plan that meets both conditions—when these conditions are interpreted quantificationally—leads to contradiction, as we've just seen.

Once we have a handle on how inconsistency-explanations are supposed to work in MCCTS, this sort of explanation makes a great deal of sense. It requires next to no exertion for someone who is familiar with parametric interpretation of modal notions. We all know that modal propositions are sensitive for their truth to—that they have, at least from a semantic point of view, as their truth-makers—parameters like accessibility relations, bodies of information, and preference orderings. Once that understanding is in place, all one needs to understand why $!\phi$ and $¡\neg\phi$ are inconsistent is to realize that these sentences require, for a preference ordering to meet the cognitive instructions they specify, logically inconsistent things of that preference ordering (and thereby of the plan from which that preference ordering is derived). The satisfaction of this condition would require a modal contradiction to be true.

On Canonicity Modal Cognitivism erred in failing to provide a theoretically adequate explanation of such facts about inconsistency; it could not deliver an intelligible explanation of inconsistency in a language containing imperatives. Expressivism failed here as well, although its error had a different source. Expressivism assumed accounting for the characteristic non-representational function of imperatives would require the domain of semantic values for imperatives to be psychological in character. In so doing, it failed to respect our canonical methodology for explaining object-language inconsistency: the derivation of a contradiction in the semantic metalanguage.

We avoided this error by assuming that the domain of semantic values for imperatives might be abstract objects—plans, in the sense of the sorts of model-theoretic objects that are normally viewed as the "truth-makers" for modal notions like requirement and allowance. This allowed us to state an explanation of object-language inconsistency that was not only canonical in form, but also intelligible. So far as inconsistency goes, the account we have stated here clearly surpasses Expressivism and Modal Cognitivism.

5.4 Consequence

Here is a remarkable feature of the resulting view. It is, for all we have said, *formally indistinguishable* from the sort of Modal Cognitivism defended by [53]. Recall that Lewis endorses the notion that the value of the interpretation function for $!\phi$ and $\Box\phi$ is the same:

$$[\![!\phi]\!] = [\![\Box\phi]\!] \quad \textbf{(Lewis' Identity)}$$

Notice that, given the following truth condition for a sentence of the form $\Box\phi$—which says that $\Box\phi$ is true at Λ_a only if the Λ_a-best possibilities are ϕ-possibilities—*we may endorse Lewis' Identity too*.

$$[\![\Box\phi]\!]^{\Lambda_a} = 1 \text{ only if } \forall w \in min_{\Lambda_a} : \{w\} \models \phi$$

Here is why. The result of λ-abstracting over the sphere parameter when considering a modal sentence $\Box\phi$ is just a function from plans into propositions—potentially, the *very same function* as is expressed by the corresponding imperative.

$$[\![\Box\phi]\!] \; = \; \lambda\Lambda_a \; . \; 1 \text{ only if } \forall w \in min_{\Lambda_a} : \{w\} \models \phi$$

Indeed, I will suggest (for reasons that will become clearer in a moment) that we endorse the identity that Lewis proposes. Whatever function is expressed by $[\![!\phi]\!]$, it is the same function that is expressed by $[\![\Box\phi]\!]$.

This is remarkable for two reasons. The first reason is that it is worrying: we are trying to get away from Modal Cognitivism. In endorsing Lewis' Identity, we are identifying the semantic value of an imperative with that of its corresponding modal sentence. The second reason is that it is, well, surprising. We began from the notion that imperatives *mean* properties of plans. And here we are contemplating the identification of the semantic value of an imperative with what is, according to Modal Cognitivism, its modal "translation". Indeed, the fact that the semantics we have proposed—even prior to our endorsement of Lewis' Identity—is even *compatible with* Lewis' Identity would seem to suggest that something has gone wrong. The semantic value of an imperative *cannot be* a proposition. So the correct semantics for imperatives cannot be compatible with the semantic value of an imperative being a proposition.

What went wrong? Actually nothing at all. The *meaning* of an imperative $!\phi$ is, I insist, a property of a plan. The *meaning* of its corresponding modal sentence $\Box\phi$ is, I insist, a proposition (to the effect that the salient plan requires ϕ). The semantic function of an imperative $!\phi$ is to partition plans into allowed (those that require ϕ) and disallowed (those that do not require ϕ) classes. The semantic function of its corresponding modal sentence $\Box\phi$ is to make a claim about whether the plan in

question does, in fact, require ϕ. In a slogan then: *imperatives partition plans; modal sentences partition logical space*.

In short, the *content* of $\Box\phi$ is not represented with an expression of the form $[\![\Box\phi]\!]$. Such an expression gives (something like) its Kaplanian *character*. Instead, its content is represented with an expression of the form $[\![\Box\phi]\!]^{\Lambda_a}$. The content of the corresponding imperative $!\phi$ is, however, given by an expression of the form $[\![!\phi]\!]$. An expression of the form $[\![!\phi]\!]^{\Lambda_a}$ (the value of which is either 0 or 1) gives its *extension* relative to Λ_a (as, to compare, the result of feeding a property an object will give that property's extension relative to that object). Sentences serve to proffer their contents at contexts of utterance. Imperatives proffer objects of an entirely different sort than their corresponding declaratives.

In fact, far from causing us trouble, taking on Lewis' Identity allows us to give a treatment of imperative consequence that might have seemed proprietary to the Modal Cognitivist. The consequence relation for an imperative language is precisely the same as the consequence relation for a corresponding modal language. Indeed we can define a general consequence relation covering both the imperative and modal language as follows.[34] Letting $\Phi_1, ..., \Phi_n, \Psi$ be sentences of any type at all:

$$\Phi_1, ..., \Phi_n \models \Psi \text{ iff } \forall\Lambda_a : [\![\Phi_1]\!]^{\Lambda_a} = ... = [\![\Phi_n]\!]^{\Lambda_a} = 1 \Rightarrow [\![\Psi]\!]^{\Lambda_a} = 1$$

Informally, according to this definition, Ψ is a consequence of $\Phi_1, ..., \Phi_n$ just if, for any Λ_a, the truth of $\Phi_1, ..., \Phi_n$ at at Λ_a implies the truth of Ψ at Λ_a. Since, for any imperative $!\phi$ and any plan Λ_a, $[\![!\phi]\!]^{\Lambda_a} = 1$ iff $[\![\Box\phi]\!]^{\Lambda_a} = 1$, it follows immediately that the logic of ! is isomorphic to the logic of $\Box$: imperatives have the same logical profiles as their corresponding modal sentences.

5.5 Compositionality and Conditionals

As I mentioned above, standard semantic treatments of the indicative conditional, though they vary considerably, are available only on the assumption that the consequent of the indicative conditional *contributes its truth condition* to the computation of the meaning of the entire conditional.

- For [72, 73] as well as [51], the indicative conditional is an operator taking the antecedent proposition (and a similarity ordering over worlds) to a universal quantifier over the worlds in the consequent proposition; the content of an indicative of the form $(if\,\phi)(\psi)$ can be glossed as *all the nearest ϕ-worlds are ψ-worlds*.
- For [46], indicative consequents express generalized quantifiers over worlds, the domain of which is restricted by the proposition expressed by the antecedent. Whatever quantificational modality M is expressed by the consequent, the logical form of a conditional of the form $(if\,\phi)(\psi)$ is given by something like $M(\phi)(\psi)$ and can be glossed as M ϕ-worlds are ψ-worlds. So, if the modal is a deontic necessity modal, the meaning of a conditional of the form $(if\,\phi)(should\psi)$ is given by something like *all the best ϕ-worlds are ψ-worlds*.

[34] I am here ignoring the complication that truth may ultimately be relative to other things than plans.

The Lewis-Stalnaker semantics seems to require the consequent to characterize a non-vacuous propositional condition on the nearest antecedent-worlds. The Kratzer semantics seems to require the consequent to express quantification over worlds—i.e., a proposition. So while our semantics has so far managed to equal or surpass Modal Cognitivism, indicative imperative conditionals (hereafter again CIs) may appear to be a stumbling block for it.

This is, in fact, not the case. It is clear enough what CIs should mean on my account: like basic imperatives, they express properties of plans. Which property is that? Clearly: a CI of the form $(if\,\phi)(!\psi)$ expresses the property of preferring ψ to $\neg\psi$ **conditional on** ϕ! In other words, a CI expresses the property a plan has when it serves as the truth-maker for the corresponding conditional obligation.

$$[\![(if\,\phi)(!\psi)]\!] \;=\; \lambda\Lambda_a\,.\,1 \text{ only if } [\![(if\,\phi)(\Box\psi)]\!]^{\Lambda_a} = 1$$

So it is relatively clear what CIs ought to mean on our account. Indeed, if we help ourselves to the standard Kratzer semantics for deontic conditionals—on which a conditional of the form $(if\ \phi)(\Box\psi)$ is true relative to Λ_a is true only if the Λ_a-best ϕ-worlds are ψ-worlds, we arrive at the following explicit proposal for CIs.[35]

$$[\![(if\,\phi)(!\psi)]\!] \;=\; \lambda\Lambda_a\,.\,1 \text{ only if } \forall w \in min_{\Lambda_a}(\|\phi\|) : \{w\} \models \psi$$

Compositional Implementation Nevertheless, to generate this meaning *compositionally*—and to identify an imperative's compositional semantic contribution when it is embedded—is another matter entirely. Suffice it to say that (i) it can rather easily be done, (ii) but doing it is a fairly technical matter. The reader who is not interested in such technicalities may skip to the next section.

Before doing this, I should say that ultimately I reject both the Stalnaker-Lewis and Kratzer analyses of indicative conditionals, instead preferring to treat indicative antecedents as context-shifters for the interpretation of their consequents (cf. [28, 36, 42, 89, 91]). Compositionally generating the correct meanings for CIs is automatic on this understanding of indicatives (see [[14], 178ff] for discussion). Since, however, I do not wish to spend the time motivating this way of thinking of indicatives here, I will here gesture at how one might compositionally implement my analysis of imperatives within the linguistically standard paradigm for thinking about indicative conditionals, namely, Kratzer's.

One idea would be to appropriate the Kratzer semantics, but to suggest that, instead of expressing generalized quantifiers (i.e., functions from propositions and propositions into {0, 1}), imperatives express functions from propositions and propositions into functions from *plans into* {0, 1}. We could do this very easily by treating the imperative operator ! and the modal operator $\Box$ as of the *same semantic type*,

[35] For a variety of reasons, I think the Kratzer semantics is inadequate (cf. [28, 91]), but to simplify I will work with it here.

but treating imperatives as having a missing argument place (represented with a variable X) that is bound by λ-abstraction.[36] For genuine modals, the argument place is filled by a plan (modals are interpreted relative to plans). For imperatives, the argument place is left unfilled. Here is how one compositional implementation of this idea might work, in an intensional Montague Grammar. The left tree gives a modal logical form, while the right gives an imperative logical form.

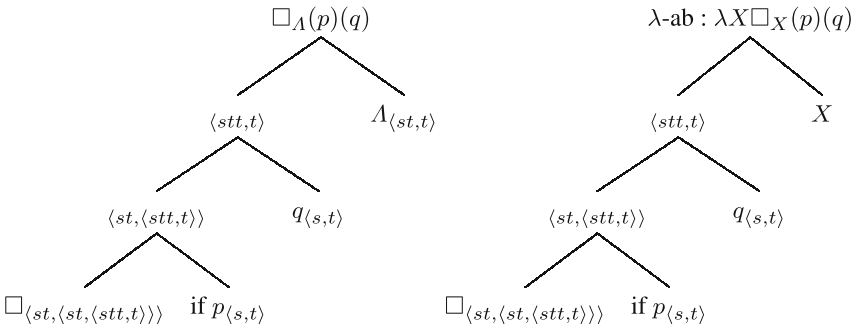

As is standard, notation of the form $\langle \tau, \tau' \rangle$ expresses a semantic type: a function from objects of type τ into objects of type τ'. When the meaning is clear, we abbreviate $\langle \tau, \tau' \rangle$ as $\tau\tau'$. Type s is the type of possible worlds, while type t is the type of verdicts, i.e., $\{0, 1\}$. Propositions are treated as of type $\langle s, t \rangle$, while plans are (simplistically) treated as of type $\langle st, t \rangle$, i.e., sets (equivalently, characteristic functions) of propositions.

5.6 Comparison

The proposal I have described in this section has a number of important antecedents. The idea that the semantics of imperatives essentially involves a kind of reference to a plan or body of preferences—less abstractly, that imperatives, as a matter of their meaning, express information about planning—can be found in (at least) [13, 14, 31, 56, 58, 59, 70, 75]. In a sense each of these authors treats imperatives as expressing or characterizing a property of an action-guiding semantic parameter. None of these theories, however, uses the hypothesis that the property in question is *modal* in nature to do any work. So, none of the explanations I have pursued in this section—all of which essentially appeal to the non-propositional, but nevertheless modal, "content" of an imperative to explain the relevant phenomena—can be replicated by these theories.

[36] If we're working in a Categorial or Variable-Free Grammar (see [38]), we could alternatively treat the open argument place as expressing the identity function $\lambda X_{\langle st,t\rangle}.X$ and generate the correct meaning for the sentence by combining this semantic value with that of its sister node by the composition operation $\lambda f \lambda g \lambda X.f(g(X))$.

This is not necessarily a strike *against* these theories—perhaps they can develop proprietary explanations of the phenomena in question. Generally, this is hard to assess, since in most cases I am not sure how the author would choose to explain the phenomena. But we can at least make a start. I will do this with a brief discussion of the theories that seem to me to be the most well-founded, developed, and compelling: Portner's and Starr's.[37] My discussion of Portner is mostly critical, while my discussion of Starr is generally supportive (with one major caveat). (The descriptions of the theories I give here are compressed, so the reader who is unfamiliar with the theories may wish to gloss over the details or skip directly to Section 6).

5.6.1 Portner's Account

The central phenomenon Portner's account tries to explain is the tendency for utterances of imperatives to establish the truth of prioritizing necessity modals. Uttering *take out the trash* at a context generally makes it the case that the addressee should or must take out the trash. Here is how he proposes to account for this phenomenon. Portner understands the semantic value of an imperative to be a property (defined only for the addressee—roughly, an action[38]), while endorsing the existence of a canonical relationship between semantic type and force. Declaratives express propositions, hence are canonically used to proffer propositions for acceptance (addition to one's beliefs, to the Common Ground, etc.). Imperatives express actions, hence are canonically used to proffer actions for performance (addition to one's To-Do List). Letting c be a context, T_c a function from agents to their To-Do Lists, α_c the addressee of c, and $c[\Phi]$ the result of updating c in accordance with an utterance Φ at c, Portner's theory of imperative meaning can be summed up with the following simple definition:

Update Your To-Do List
When $[\![\Phi]\!]^c$ is an action, $T_{c[\Phi]} = T_c^{\alpha_c \to T_c(\alpha_c) \cup \{[\![\Phi]\!]^c\}}$

The To-Do List is subsequently used to define a preference ordering on alternatives (which is subsequently used in the interpretation of modals).

Portner's account has several difficulties which I have discussed elsewhere [[14], Chs. 2-3]. In spite of appearances, it is an Expressivist account; the meaning of an imperative, for Portner, *consists in* its To-Do List-modifying function: an account of how imperatives tend to establish the truth of corresponding prioritizing necessity modals explains "everything that needs to be explained about [their] meaning" [[59], 366]. (It is thus a bit confusing to say that the semantic value of an imperative is an action, since, for Portner, the semantic value of an imperative plays a rather subsidiary role in his theory of imperative meaning.) As such, Portner's account is saddled with

[37] For strong criticisms of Han and Mastop, see, respectively, [60] and [75].
[38] Barker [5] concurs that imperatives express actions, although his broader theory is quite different from Portner's.

the problems that tend to saddle Expressivist accounts. For instance, he lacks an account of imperative inconsistency.

Portner's account does not suggest an obvious analysis of conditional imperatives (and embedded imperatives more generally). It will not do to treat conditional imperatives (CIs) as adding a conditional proposition or action to an addressee's To-Do List. This amounts, in effect, to treating CIs as having wide-scope logical forms, i.e., to representing a CI with an expression of the form $!(\phi \rightarrow \psi)$. Such a proposal is inadequate for reasons we saw in Section 2.3. Nor will it do to say that CIs conditionally update the addressee's To-Do List (conditional on the antecedent's being accepted or true at the context of utterance); CIs establish the truth of conditionalized prioritizing necessity modals, even (indeed, especially) when their antecedents are not accepted. An utterance of *if you rob Jones, do it gently* tends to make it the case that: if you rob Jones you should (or must) do it gently. It does so even when the addressee has no intention to rob Jones; its intuitive function is to constrain Jones' behavior in certain contingencies that, even if not actual, might later become actual. An account on which the effect of a CI is *conditioned* cannot explain this rather central fact about their meaning.[39]

Finally, the property of To-Do Lists that Portner has an imperative characterizing—the property of having the prejacent on the To-Do List—is not theoretically adequate. A To-Do List can permit $\neg\phi$ even when ϕ is on the To-Do List; the presence of an action in an ordering source does not guarantee that every world that is maximally good with respect to the resulting ordering is one where that action is done.[40] Things in the ordering source are ceteris paribus preferred, but in cases where cetera are not paria, they can fail to be preferred (thus permitting their absence). But imperatives are *always and everywhere logically inconsistent* with contrary permissions. The notion of a To-Do List that both accepts an imperative and validates its contrary permission sounds to me suspiciously close to a contradiction. But it is something to which Portner's account seems to be committed. To sum up,

[39]This isn't to say there is no way to model the update that CIs might "execute" on a context; for such a model see [[14], 178ff]. But adopting such a model would require a rather major reworking of Portner's theory—one that is not obviously available to him. Why? The idea that the semantic value of an unembedded imperative is an action fails to generate a plausible semantic value for a CI: saying that a CI's semantic value is typed as an action collapses the theory into wide-scope-ism (since Portner's theory would have this action be added to the To-Do List), while saying it is typed as something else is, insofar as (i) we expect clauses of the same type to have semantic values of the same type (see [58]) and (ii) conditionals are, per [7], typed as clauses according to the type of their matrix, rather than subordinate, clause (so that CIs are, in fact, typed *as imperatives*), prima facie undesirable.

This said, Modal Noncognitivism in no way requires an account of how CIs update the context, contra Portner's [60] remark that an "issue for [Modal Noncognitivism] is that it does not define a particular update to the context for a given imperative." Modal Noncognitivism's theory of imperative meaning is given, not by a theory of how imperatives update the context, instead by the semantic identities proposed in this section.

[40]This is something of which Portner is aware—indeed, enthusiastic. See his [61], which uses this fact in an analysis of permission interpretations of imperatives (e.g., *go ahead, have an apple!*). I have argued elsewhere that this is not a satisfactory treatment of permission imperatives [[14], 15ff]. Even if it were, this feature of the analysis would be undesirable in light of the argument I am giving here.

the motivating idea of Portner's account (one shared by Han's, Mastop's, and Starr's accounts)—that imperatives, as a matter of their meaning, express information about planning—is spot-on. But Portner's implementation of that idea seems to me flawed: he cannot, in the end, give satisfactory accounts of the phenomena outlined in Section 2, all of which are (or should be regarded as) central to an adequate theory of imperative meaning. Modal Noncognitivism can.

5.6.2 Starr's Account

Starr [75] states what he calls a "preference semantics" for imperatives. For Starr, giving an imperative semantics amounts to giving an account of how imperatives update preference orderings; specifically, an imperative of the form $!\phi$ updates a preference state so as to establish a preference for ϕ to $\neg\phi$. The apparatus Starr develops to cash out this idea is impressive, and I think it's best to refer the reader interested in the details to his discussion. Setting the apparatus to the side, however, this is a very nice idea, elegantly implemented, with the potential to account for an extraordinarily large range of phenomena. Although our accounts developed independently, I would guess that we agree about the explanations of most of the phenomena in question.

Because Starr's account is couched in a "dynamic" semantics—i.e., a semantics which assigns update functions to sentences (see, e.g., [82])—he has a ready account of CIs. Informally speaking, conditionals are interpreted, in dynamic semantics, as context-shifters for the execution of the update expressed by their consequent. So, on the dynamic treatment (unlike in the theories of Kratzer, Stalnaker, and Lewis), conditionals are promiscuous with respect to the semantic type of (i.e., type of update expressed by) their consequents. I have no objection to this idea for CIs—indeed, there is reason to think it is roughly as explanatory as the account of CIs I have suggested here, and I have outlined such an account elsewhere [[14], 178ff].

Starr's probable account of imperative inconsistency can be inferred from his notion of "preferential consequence." He writes: "Preferential consequence...tracks the preservation of which alternatives are best... ψ is a consequence of some premises just in case accepting ψ after accepting the premises never changes which alternatives are best" (34). Preferential inconsistency of a set of sentences probably, then, will amount to the *non-existence* of a preference state that accepts these sentences in sequence. This is roughly the same notion of inconsistency as the one we have deployed: recall that the inconsistency of an imperative of the form $!\phi$ and a contrary permission of the form $¡\neg\phi$ was accounted for in Section 5.3 by pointing out that the assumption that there was some plan Λ_a such that $\Lambda_a \in [\![!\phi]\!]$ and $\Lambda_a \in [\![¡\neg\phi]\!]$ (i.e., a plan that accepts both sentences) generated a contradiction.

There is one point of difference that may matter for purely (merely?) philosophical reasons. Why, for Starr, *is there no* preference state that accepts, e.g., both $!\phi$ and $¡\neg\phi$? There are two possibilities: (i) the existence of such a state would entail a contradiction, or (ii) there are barely possible (i.e., non-contradictory) states that accept both $!\phi$ and $¡\neg\phi$, but such states are banished, by stipulation, from the semantics, due, perhaps, to their manifest irrationality. The latter strategy is, I think, unpalatable: it would seem to foreclose to Starr the canonical way of accounting for object

language inconsistency, i.e., the derivation of a contradiction in the metalanguage.[41] In any case, he does not opt for it.

Starr's adoption of dynamic semantics makes taking the former strategy more difficult than it otherwise would be.[42] Pursuing such a strategy requires an account of (i) the condition imperatives enforce on a preference state, (ii) the condition under which a preference state accepts a permission such that imperatives and contrary permissions place *logically incompatible conditions* on preference states. Imperatives, in all likelihood, would have to make their prejacents *required by* the target preference states. But the change a preference state undergoes when it shifts to requiring ϕ is *very* different from the simple shift involved in simply adding ϕ to your stock of preferences. The former requires an account of how prior preferences—ones that conflict with the preference for ϕ so as to block the global preference state from requiring ϕ—are *revised* so that conflict is ameliorated. Such an account would appear to be a matter of proper concern for epistemology, rather than semantics; we should not build a substantive theory of (rational) preference-revision into any part of our semantics.[43]

I say this as a proponent of a picture of semantics on which a semantic theory for a sentence is a theory about what sort of cognitive instruction that sentence proffers (Section 5.1). However, cognitive instructions, as I understand them, are relatively *thin* things. An assertion that ϕ might accurately be described as a cognitive instruction whose force is to propose to get the addressee to accept that ϕ. When ϕ conflicts with an agent's prior information, then, in order to fully accept the assertion, the agent cannot simply add the proposition that ϕ to her stock of beliefs. She must revise her prior beliefs to eliminate the source of conflict. Not just any old revision will do; the revision must be, for example, conservative with respect to those priors. How rational agents revise priors conservatively is the topic of a vast amount of work in formal and informal *epistemology*, rather than natural language semantics.

That seems to me like the correct division of theoretical labor. The job of natural language semantics is to associate sentences with general, and relatively skeletal, cognitive instructions. These have a plausible claim to being what is meant by utterances in natural language. It is the job of epistemologists (and theorists interested in

[41] For related reasons, we have reason to prefer a strategy which treats states that accept, for example, $\Phi \wedge \Psi$ without accepting Φ as contradiction-licensing (rather than barely possible, but irrational).

[42] What follows is not to be understood as a critique of dynamic semantics generally (except on the strongest understanding of the dynamic semantic program, on which the semantic value of any sentence is a deterministic program or update function). The dynamic semantic approach to, to pick an example at random, donkey anaphora (see, e.g., [30]) plausibly does not suffer from the problems I describe here.

[43] Here is some detail about Starr's account to fill this point in. A *preference state* R is defined as a binary relation on a set of alternative propositions (typed as sets of worlds) such that $\langle a, a' \rangle \in R$ iff *a is preferred to* a'. For Starr, imperatives update preference states by strengthening their constituent preferences: updating R with $!\phi$ yields $R[!\phi]$ such that (i) $R \subseteq R[!\phi]$, (ii) $R[!\phi]$ contains a preference for ϕ to $\neg\phi$, i.e., $\langle ||\phi||, ||\neg\phi|| \rangle \in R[!\phi]$, and (iii) if $\langle a, a' \rangle \in R$, $\langle a[\phi], a - a[\phi] \rangle \in R[!\phi]$ (where $a[\phi] = a \cap ||\phi||$, i.e., is the result of intersecting a with the proposition that ϕ). But component (iii) of the update amounts to a view about how prior preferences are revised on the adoption of a new preference (one that, in all likelihood, would need to be significantly complicated to yield a realistic account of its target). For a sense of the complexities involved here, see [81].

rational changes in attitudes, more generally) to give a substantive account of how rational agents can comply with these instructions. It is implausible to think that the meanings of our utterances would build in anything like a substantive account of rational compliance. (For one thing, people can disagree about what rational compliance consists in without thereby disagreeing about what their utterances mean.) For this reason, I think it would be a mistake to augment the theory of assertion with the fruits of the epistemological literature on belief-revision. For similar reasons, I think it is a mistake to augment an account of the cognitive instruction proffered by an imperative with a substantive theory of (rational) preference-revision. Insofar as a semantics for imperatives commits itself to such a theory—as the dynamic component of Starr's semantics for imperatives, in fact, seems to do—it is inherently implausible.

In a slogan: semantics furnishes a theory of cognitive directives. Epistemology furnishes a substantive theory of diachronic compliance for cognitive directives. What is left to the semantics? It would seem to be just the idea that imperatives propose to make their prejacents required. This is awfully close to Modal Noncognitivism.[44]

6 Situating Ourselves

Modal Noncognitivism is an account of considerable power. It retains the "virtues" of Modal Cognitivism—indeed, makes these merely apparent virtues actual—while avoiding the pitfalls that arise from the identification of the meaning of an imperative with a proposition. To conclude this paper, I want to expand on these remarks, as well as reflect on certain philosophical and methodological issues the approach defended here might raise.

6.1 On Function and Proposition

Imperative and permissions have properties of plans, rather than propositions, as their semantic values. Such objects are of exactly the right sort to explain the non-representational use that imperatives and permissions characteristically have. An imperative $!\phi$ semantically specifies a property than an agent can try to *psychologically approximate* (the property of having a plan according to which ϕ is required, just as an ordinary declarative specifies a property that an agent can try to *psychologically approximate* (the property of having beliefs such that the situations compatible with what the agent believes all realize the sentence). An agent approximates the

[44]The main remaining difference is that Starr does not suggest an explicitly modal understanding of notions like requirement and permission. Starr does aim to give accounts of phenomena we've decided to set to the side, e.g., Ross' Paradox, but we do not, I think, disagree about the explanation of these phenomena. (For my treatment of Ross' Paradox, see [[14], 137ff].) For this reason, I would disagree with his claim that "The preference semantics for imperatives developed above enjoys important empirical and conceptual advantages over existing accounts."

property semantically specified by $!\phi$ by adjusting her plans so that they could be represented as being such that ϕ is required with respect to them.

As the above paragraph suggests, the suggested connection between the semantics and characteristic function of imperatives is an instance of a more general connection between semantics and characteristic function. I have suggested that, as a general matter, sentences semantically function to resolve cognitive issues, and, therefore, that sentences' semantic values encode cognitive instructions. A sentence encodes such a cognitive function by semantically expressing a property of (an abstract representation of) a cognitive state. In the case of an ordinary declarative, such a property is straightforwardly derivable from its ordinary (propositional) semantic value.

6.2 On Function and Semantic Type

There is some tension between the idea that the semantic value of a declarative is a proposition, whereas the semantic value of an imperative is a characteristic function of plans. What, say, is the semantic value of a sentence of the form $(!\phi \wedge \psi)$ or $(\phi \wedge !\psi)$? (The examples are due to [75].)

(47) Go home and I'll go to the grocery store
(48) I'll go to the grocery store and you go home

Of course, following the dictum that imperative operators take obligatorily wide scope, our toy imperative language does not allow the formation of such conjunctions. The possibility of sentences might, however, seem like a good reason to abandon the dictum.[45]

What sort of semantic value could such conjunctions have? How does one conjoin a proposition with a property of a plan? Such a worry is easily handled by simply shifting to a unified treatment of declaratives and imperatives. This may be done simply by extending our earlier treatment of the 'holds at' relation to cover imperatives. Letting a **representor** for an agent a be a pair $\langle S_a, \Lambda_a \rangle$ of a set of possible worlds S_a (representing a's information) and a set of action descriptors Λ_a (representing a's plans), here is what I have in mind. (As before, capital Greek letters stand for sen-

[45] Perhaps such examples feel forced or contrived. For more persuasive examples; see [[14], 133ff]. There are also, it should be mentioned, cases like:

(49) Stop or I'll shoot!
(50) Drink another beer and we'll win the competition! ([64])

Following roughly [19, 83], I claim that such sentences plausibly function to express assertions of conditionals:

(51) ≈If you don't stop, I'll shoot [perhaps conjoined to an imperative: *stop!*]
(52) ≈If you drink another beer, we'll win the competition [perhaps joined to an imperative: *drink another beer!*]

While there is a mystery about how to generate such readings, I think there is very little pressure for a semantics of imperatives to account for them. Instead, the pressure is to come up with an analysis of the relevant connectives ('or' and 'and') that allows them to *transform imperatives into subordinate clauses* functioning to restrict the domain of quantificational modals like 'will'.

tences of arbitrary type, regular Greek letters stand for sentences of the first-order base language).

Holding At, In General

$\langle S_a, \Lambda_a \rangle \models p$ iff $\forall w \in S_a : p \in w$
$\langle S_a, \Lambda_a \rangle \models \neg\phi$ iff $\forall w \in S_a : \{w\} \not\models \phi$
$\langle S_a, \Lambda_a \rangle \models [\![\phi]\!][!\psi]$ only if $\forall w \in min_{\Lambda_a}(S_a \cap \|\phi\|) : \{w\} \models \psi$
$\langle S_a, \Lambda_a \rangle \models [\![\phi]\!][¡\psi]$ only if $\exists w \in min_{\Lambda_a}(S_a \cap \|\phi\|) : \{w\} \models \psi$
$\langle S_a, \Lambda_a \rangle \models (\Phi \wedge \Psi)$ iff $\langle S_a, \Lambda_a \rangle \models \Phi$ and $\langle S_a, \Lambda_a \rangle \models \Psi$

Some things to notice about this definition. The first is that CIs are treated as basic; regular imperatives—imperatives of the form $!\phi$—can be represented with an expression of the form $[\top][!\phi]$ (where $\top$ is any tautology). A regular imperative of the form $!\phi$ thus holds at a representor $\langle S_a, \Lambda_a \rangle$ just if $\forall w \in min_{\Lambda_a}(S_a) : \{w\} \models \phi$—just if, that is to say, the best worlds compatible with a's information, given a's plans, are ϕ-worlds. (All this goes, mutatis mutandis, for indicative conditional permissions).

Second, and more interesting, there is no problem understanding a conjunction of the form $(!\phi \wedge \psi)$. This expresses the property a representor has iff it satisfies the property expressed by $!\phi$ as well as that expressed by ψ. Effectively, then, we are claiming that sentences of any type express the very same sort of semantic object: in general, *sentences express properties of representors*. Imperatives express properties that bear on the planning parameter of a representor, while declaratives express properties that bear on its informational parameter. Sentences are not typed according to the semantic type of their denotation, since all sentences express characteristic functions of representors. Instead, they are typed according to the kind of property of a representor they express.

Pending further discussion below, there is, I want to suggest, nothing very radical about this idea. The property a declarative expresses of a representor is, as I have already explained, derivable from the standard, propositional semantics for declaratives. And even though declaratives and imperatives express semantic values of homogenous semantic type, we are still able to give a very intuitive account of why (and how) the canonical function of a declarative differs from the canonical function of an imperative. Declaratives express properties that bear on an information-tracking parameter, and this is a property than an agent can try to psychologically approximate (by modifying her beliefs so that they are representable with an information-tracking parameter satisfying the property in question). Imperatives express properties that bear on a planning parameter, and this is a property than an agent can try to psychologically approximate (by modifying her plans so that they are representable with plan satisfying the property in question). In other words, we explain a sentence's canonical function in terms of the *content of the property* that sentence, as a matter of its semantics, expresses, rather than the gross ontological category of its semantic value.

6.3 On the Role of Propositions

Our account of imperatives was designed with an eye to retaining the advantages of Modal Cognitivism. It did this by, in a sense, assigning imperatives truth conditions

(or, more precisely, truth-of-a-plan conditions). An imperative is *true of* a plan Λ_a just if its "corresponding modal sentence" is *true relative to* Λ_a. Similarly, the predicate P is true of an object o just if the sentence $P(x)$ is true relative to (an assignment mapping x to) o, and the quantifier $\mathcal{Q}$ is true of a property $\mathcal{P}$ just if the sentence $\mathcal{Q}(X)$ is true relative to (an assignment mapping X to) $\mathcal{P}$.

And indeed the character of explanation for the semantic behavior of such things (particularly quantifiers) looks very much like the character of our explanation of the semantic behavior of imperatives. I have suggested treating imperatives, semantically, as modal quantifiers *minus an argument*. Analogously, consider the case of quantifiers missing, e.g., a scope. Such quantifiers stand in semantic relationships to one another; for instance $\forall$ entails $\exists$ and is inconsistent with $\exists\neg$ relative to a single resolution of the missing scope. An explanation of this fact does not cease to be truth-conditional simply because it makes no sense to say that $\forall$ or $\exists$ or $\exists\neg$ is true or false simpliciter.

Something like this is also the case for imperatives. There is a sense in which our explanations of Facts 1, 2, and 3 ultimately bottom out in the truth-conditional characteristics of modal propositions, i.e., a sense in which the relevant explanations are fundamentally truth-conditional in character. For instance, our explanation for why an imperative of the form $!\phi$ is inconsistent with a permission of the form $¡\neg\phi$ trades on the fact that a state of affairs in which ϕ is required and $\neg\phi$ permitted is a *contradictory state of affairs*. So it is, I think, substantially correct to say that our analysis of imperatives and permissions assigns these sentences truth conditions, by way of accounting for their semantics.

But are imperatives true or false? More precisely, does it ever make sense to ascribe truth or falsity simpliciter to an imperative? It does not. Of course, such ascriptions are not grammatical. Deeper than this, however, there is the fact that, imperatives express the wrong kinds of semantic values to be properly called true or false. They do not encode a picture of the world *as-is*, what we might term, following the notion of [54] that coming to believe something is a matter of self-locating in a space of possibilities, a 'locational' perspective. They encode a picture of the world *to-be*, what we might term a 'directional' perspective. Despite having truth conditions, imperatives are neither true nor false.

I will suppose that this is a congenial way of resolving the issue of whether imperatives have truth conditions. There is still, however, a worry in the neighborhood having to do with declaratives. It is hard to see in what sense declaratives might be (or be called) true or false *simpliciter*.[46] At a certain level of abstraction, declaratives lack truth conditions in the same way that imperatives do: they express objects (namely, properties) that can only be evaluated for truth or falsity *relative to a representor*.

In reply, I will suggest that sentences can be evaluated as true or false *simpliciter* insofar as *they encode a perspective that an agent can evaluate for truth or falsity*. Only certain kinds of perspective are evaluable for truth or falsity, namely, those that encode a picture of a way the world might be—only, that is to say, locational perspectives. I have no argument for this beyond its inherent plausibility: it is possible

[46] I am grateful to Andrew Alwood for pressing this worry.

to evaluate a locational perspective (crudely, a belief) for truth, it is not possible to evaluate a directional perspective (crudely, a plan) for truth.[47]

But I can at least say something about how I understand the notion of evaluating a sentence for truth or falsity. I will suggest that evaluating a sentence for truth, in the sense in which I am interested here, involves executing a test on one's cognitive state. In particular, one tests for a sentence's truth by testing [in the sense of [82]] whether it is accepted relative to one's cognitive state. The idea is something like this: to decide that a sentence is true is to successfully test one's cognitive state for its acceptance; to decide it is false is to successfully test for its negation. Here is a quick effort at formalizing this notion. Let us first introduce truth and falsehood predicates $\top$ and $\bot$ into our language. Writing the result of executing the test expressed by a sentence Ψ on a state $\langle S_a, \Lambda_a \rangle$ as $\langle S_a, \Lambda_a \rangle[\Psi]$, and assuming that $||\cdot||$ is *defined only* for sentences encoding locational perspectives, here is what I have in mind.

Evaluating for Truth/Falsehood

$$\langle S_a, \Lambda_a \rangle[\top(\Phi)] = \begin{cases} \langle S_a, \Lambda_a \rangle \text{ if } S_a \subseteq ||\Phi|| \\ \emptyset \text{ if } S_a \nsubseteq ||\Phi|| \\ \text{undefined otherwise} \end{cases}$$

$$\langle S_a, \Lambda_a \rangle[\bot(\Phi)] = \begin{cases} \langle S_a, \Lambda_a \rangle \text{ if } S_a \subseteq ||\Phi||' \\ \emptyset \text{ if } S_a \nsubseteq ||\Phi||' \\ \text{undefined otherwise} \end{cases}$$

There are two ways for a test expressed by a sentence to result in failure. One, the test can *fail to make sense* (as is the case here when Φ fails to encode a locational perspective, in which case '$||\Phi||$' fails to refer, and hence neither $S_a \subseteq ||\Phi||$ nor $S_a \nsubseteq ||\Phi||$). The result of testing in such cases is undefined, since no coherent test is actually expressed; in such cases, the sentence in question is non-evaluable for truth. Two, the test can make sense—in which case the sentence is at least evaluable for truth—but *be failed*; in this case, the sentence $\top(\Phi)$ is simply rejected. This is not, notice, to say that $\bot(\Phi)$ is accepted; the result of executing the test (the absurd state, $\emptyset$) represents the unacceptability of a determinate verdict of truth—something that can obtain in cases of uncertainty about Φ—rather than the acceptability of a verdict of falsehood.

The suggested understanding of evaluating for truth and falsehood is deflationary in one sense (since a sentence being evaluated as true at a state amounts to nothing more than its acceptance at that state), and inflationary in another (since only sentences that express honest-to-goodness propositions—that encode honest-to-goodness locational perspectives—will be evaluable for truth). If an indicative conditional, e.g., does not express a proposition (cf. [1, 52])—if it tells an agent, not how to self-locate, but instead, perhaps, to assign a conditional probability of 1 to the consequent on the antecedent—it is not evaluable for truth in the above sense (although it is, of course, evaluable for truth in *some* sense).

[47] One way to appreciate at least the second claim here is by considering Dreier's [20] Hiyo example. In a language where the sentence 'Bob is hiyo' expresses, by stipulation, the speect act of accosting Bob, this sentence, despite being syntactically declarative, is not evaluable for truth.

6.4 The End

It is obvious, as I have already noted, that adopting the sort of strategy I am advocating here is revisionary with respect to at least some of what counts as standard practice in semantic theorizing. This is due, at least in part, to its implicit rejection of explanations of, e.g., the differences in behavior between declaratives and other sentence-kinds that appeal to the semantic value of the declarative being *a proposition*. On this strategy, the semantic value of *any* sentence is a property; no sentences, strictly speaking, express propositions as their semantic values.

Nevertheless, in certain cases—indeed, I would claim, just those cases where truth-conditional theorizing has witnessed its greatest successes—propositions play a fairly ordinary explanatory role. In spite of all that we have said, declaratives remain truth-evaluable *because* they have propositional "contents" (although not *qua* semantic values)—because the perspectives they encode are equivalent to representations of various ways the world might be.

Imperatives, on the other hand, are not truth-evaluable because they lack propositional content; they say where to go, rather than where you are. While accommodating this platitude motivates reworking of some of the architecture of semantic theorizing, we have seen that it does not require that we start from scratch. That is a congenial result. More importantly, it is one on which future work can readily build.

References

1. Adams, E.W. (1975). *The logic of conditionals: an application of probability to deductive logic*. Dordrecht: Reidel.
2. Aloni, M. (2007). Free choice, modals, and imperatives. *Natural Language Semantics*, *15*, 65–94. doi:10.1007/s11050-007-9010-2.
3. Alwood, A. (2010). Imperative clauses and the Frege-Geach problem. *Analysis*, *70*, 105–117. doi:10.1093/analys/anp146.
4. Åqvist, L. (1964). Interpretations of deontic logic. *Mind*, *73*, 246–253. doi:10.1093/mind/LXXIII.290.246.
5. Barker, C. (2012). Imperatives denote actions. In A. A. Guevara, A. Chernilovskaya, R. Nouwen (Eds.), *Proceedings of Sinn und Bedeutung 16* (Vol. 1, pp. 57–70).
6. Belnap, N.D. Jr., & Perloff, M. (1988). Seeing to it that: a canonical form for agentives. *Theoria*, *54*, 175–199.
7. Bhatt, R., & Pancheva, R. (2006). Conditionals. In M. Everaert & H. van Riemsdijk (Eds.), *The 1576 Syntax companion: volume I*. Oxford: Blackwell. http://www-rcf.usc.edu/~pancheva/bhatt-pancheva_syncom.pdf.
8. Blackburn, S. (1984). *Spreading the word*. Oxford: Clarendon Press.
9. Blackburn, S. (1998). *Ruling passions*. Oxford: Oxford University Press.
10. Cariani, F. (to appear). Ought' and resolution semantics. *Noûs*. doi:10.1111/j.1468-0068.2011.00839.x.
11. Carr, J. (2012). Subjective ought. Ms. http://web.mit.edu/jrcarr/www/pdfs/subjectiveought.pdf.
12. Castañeda, H.N. (1960). Imperative reasonings. *Philosophy and Phenomenological Research*, *21*, 21–49. doi:10.2307/2104787.
13. Charlow, N. (2010). Restricting and embedding imperatives. In M. Aloni, H. Bastiaanse, T. de Jager, K. Schulz (Eds.), *Logic, language, and meaning: selected papers from the 17th Amsterdam Colloquium, ILLC*. http://www.springer.com/computer/theoretical+computer+science/book/978-3-642-14286-4.
14. Charlow, N. (2011). *Practical language: its meaning and use*. Ph.D. Dissertation, University of Michigan. http://www.natecharlow.com/work/dissertation.pdf.

15. Charlow, N. (to appear a). The problem with the Frege-Geach problem. *Philosophical Studies*. doi:10.1007/s11098-013-0119-5.
16. Charlow, N. (to appear b). What we know and what to do. *Synthese*. doi:10.1007/s11229-011-9974-9.
17. Chellas, B.F. (1971). Imperatives. *Theoria*, *37*, 114–129. doi:10.1111/j.1755-2567.1971.tb00064.x.
18. Copp, D. (2001). Realist-expressivism: a neglected option for moral realism. *Social Philosophy and Policy*, *18*(1–43). doi:10.1017/S0265052500002880.
19. Culicover, P.W., & Jackendoff, R. (1997). Semantic subordination despite syntactic coordination. *Linguistic Inquiry*, *28*, 195–217. http://www.jstor.org/stable/4178974.
20. Dreier, J. (1996). Expressivist embeddings and minimalist truth. *Philosophical Studies*, *83*, 29–51. doi:10.1007/BF00372434.
21. Dreier, J. (2006). Negation for expressivists: a collection of problems with a suggestion for their solution. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 1). Oxford: Oxford University Press.
22. Dummett, M. (1973). *Frege: Philosophy of language*. London: Duckworth.
23. Fox, C. (2012). Imperatives: a judgemental analysis. *Studia Logica*, *100*, 879–905. doi:10.1007/s11225-012-9424-9.
24. Geach, P.T. (1958). Imperative and deontic logic. *Analysis*, *18*, 49–56. doi:10.2307/3326785.
25. Geach, P.T. (1965). Assertion. *The Philosophical Review*, *74*, 449–465. http://www.jstor.org/pss/2183123.
26. Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge: Harvard University Press.
27. Gibbard, A. (2003). *Thinking how to live*. Cambridge: Harvard University Press.
28. Gillies, A.S. (2010). Iffiness. *Semantics and Pragmatics*, *3*, 1–42. doi:10.3765/sp.3.4.
29. Green, M.S. (1997). The logic of imperatives. In E. Craig (Ed.), *The Routledge Encyclopedia of Philosophy* (pp. 717–721). New York: Routledge. http://people.virginia.edu/~msg6m/implogic.html.
30. Groenendijk, J., & Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, *14*, 39–100. doi:10.1007/BF00628304.
31. Han, C.-h. (1998). *The structure and interpretation of imperatives: mood and force in universal grammar*. Ph.D. Dissertation, University of Pennsylvania. http://www.sfu.ca/~chunghye/papers/dissertation.pdf.
32. Hare, R.M. (1949). Imperative sentences. *Mind*, *58*, 21–39. doi:10.1093/mind/LVIII.229.21.
33. Hare, R.M. (1952). *The language of morals*. Oxford: Oxford University Press.
34. Hare, R.M. (1967). Some alleged differences between imperatives and indicatives. *Mind*, *76*, 309–326. doi:10.1093/mind/LXXVI.303.309.
35. Harrison, J. (1991). Deontic logic and imperative logic. In P. Geach (Ed.) *Logic and Ethics*. Dordrecht: Kluwer. doi:10.1007/978-94-011-3352-4n_6.
36. Isaacs, J., & Rawlins, K. (2008). Conditional questions. *Journal of Semantics*, *25*, 269–319. doi:10.1093/jos/ffn003.
37. Jackson, F., & Pettit, P. (1998). A problem for expressivism. *Analysis*, *58*, 239–251. doi:10.1111/1467-8284.00128.
38. Jacobson, P. (1999). Towards a variable-free semantics. *Linguistics and Philosophy*, *22*, 117–185. doi:10.1023/A:1005464228727.
39. Jørgensen, J. (1937). Imperatives and logic. *Erkenntnis*, *7*(8), 288–296.
40. Kaufmann, S., & Schwager, M. (2009). A unified analysis of conditional imperatives. In E. Cormany, S. Ito, D. Lutz (Eds.), *Proceedings of SALT 19* (pp. 239–256). Ithaca: CLC Publications. www.semanticsarchive.net/Archive/TYzZWYyM/.
41. Kenny, A.J. (1966). Practical inference. *Analysis*, *26*, 65–75. doi:10.2307/3326284.
42. Kolodny, N., & MacFarlane, J. (2010). Ifs and oughts. *Journal of Philosophy*, *107*, 115–143. http://johnmacfarlane.net/ifs-and-oughts.pdf.
43. Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, *1*, 337–355. doi:10.1007/BF00353453.
44. Kratzer, A. (1981). The notional category of modality. In H. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts* (pp. 38–74). Berlin: De Gruyter.
45. Kratzer, A. (1991). Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: an international handbook of contemporary research* (pp. 639–651). Berlin: De Gruyter.
46. Kratzer, A. (1991). Conditionals. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: an international handbook of contemporary research* (pp. 651–656). Berlin: De Gruyter.
47. Lascarides, A., & Asher, N. (2004). Imperatives in dialogue. In P. Kuehnlein, H. Rieser, H. Zeevat (Eds.), *The semantics and pragmatics of dialogue for the new millennium*. Philadelphia: Benjamins. http://homepages.inf.ed.ac.uk/alex/pubs/imps.html.

48. Lemmon, E.J. (1965). Deontic logic and the logic of imperatives. *Logique et Analyse*, *8*, 39–71.
49. Lenman, J. (2003). Disciplined syntacticism and moral expressivism. *Philosophy and Phenomenological Research*, *66*, 32–57. doi:10.1111/j.1933-1592.2003.tb00242.x.
50. Lewis, D. (1970). General semantics. *Synthese*, *22*, 18–67. doi:10.1007/BF00413598.
51. Lewis, D. (1973). *Counterfactuals*. Malden: Basil Blackwell.
52. Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review*, *85*, 297–315. http://www.jstor.org/pss/2184045.
53. Lewis, D. (1979). A problem about permission. In E. Saarinen, R. Hilpinen, I. Niiniluoto, M.B. Hintikka (Eds.), *Essays in Honour of Jaakko Hintikka*. Dordrecht: D. Reidel.
54. Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, *8*, 339–359. doi:10.1007/BF00258436.
55. Lewis, D. (1979). Attitudes de dicto and de se. *The Philosophical Review*, *88*, 513–543.
56. Mastop, R. (2005). *What can you do?* Ph.D. Dissertation, ILLC.
57. Parsons, J. (2012). Cognitivism about imperatives. *Analysis*, *72*, 49–54. doi:10.1093/analys/anr132.
58. Portner, P. (2004). The semantics of imperatives within a theory of clause types. In K. Watanabe & R. Young (Eds.), *Proceedings of SALT 14*. Ithaca: CLC Publications. http://semanticsarchive.net/Archive/mJlZGQ4N/.
59. Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, *15*, 351–383. doi:10.1007/s11050-007-9022-y.
60. Portner, P. (to appear a). Imperatives. In M. Aloni & P. Dekker (Eds.), *Cambridge handbook of semantics*. Cambridge: Cambridge University Press. http://www.semanticsarchive.net/Archive/Tc3YTc5M/Portner-Imperatives-draftFeb152012.pdf.
61. Portner, P. (2012). Permission and choice. In G. Grewendorf & T. Zimmermann (Eds.), *Discourse and grammar: from sentence types to lexical categories. Studies in generative grammar*. Berlin: Mouton de Gruyter. http://semanticsarchive.net/Archive/jI1YjMyY/.
62. Rescher, N., & Robison, J. (1964). Can one infer commands from commands? *Analysis*, *24*, 176–9. doi:10.1093/analys/24.5.176.
63. Ross, A. (1941). Imperatives and logic. *Theoria*, *7*, 53–71. doi:10.1111/j.1755-2567.1941.tb00034.x.
64. Russell, B. (2007). Imperatives in conditional conjunction. *Natural Language Semantics*, *15*, 131–166. doi:10.1007/s11050-007-9012-0.
65. Schroeder, M. (2008). *Being for*. Oxford: Oxford University Press.
66. Schroeder, M. (2008). How expressivists can and should solve their problems with negation, *42*, 573–599. doi:10.1111/j.1468-0068.2008.00693.x.
67. Schwager, M. (2006). Conditionalized imperatives. In M. Gibson & J. Howell (Eds.), *Proceedings of SALT 16*. Ithaca: CLC Publications. http://user.uni-frankfurt.de/~scheiner/papers/schwagerFEB07.pdf.
68. Schwager, M. (2006). *Interpreting imperatives*. Ph.D. Dissertation, University of Frankfurt.
69. Searle, J.R. (1962). Meaning and speech acts. *The Philosophical Review*, *71*, 423–432. http://www.jstor.org/pss/2183455.
70. Segerberg, K. (1990). Validity and satisfaction in imperative logic. *Notre Dame Journal of Formal Logic*, *31*, 203–221. doi:10.1305/ndjfl/1093635415.
71. Segerberg, K. (1992). Getting started: beginnings in the logic of action. *Studia Logica*, *51*, 347–378. doi:10.1007/BF01028968.
72. Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory (American Philosophical Quarterly Monograph Series 2)* (pp. 98–112). Oxford: Blackwell.
73. Stalnaker, R. (1975). Indicative conditionals. *Philosophia*, *5*, 269–286. doi:10.1007/BF02379021.
74. Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and semantics 9: Pragmatics*. New York: Academic Press.
75. Starr, W.B. (2011). A preference semantics for imperatives. Ms.
76. Stenius, E. (1967). Mood and language game. *Synthese*, *17*, 254–274. doi:10.1007/BF00485030.
77. Swanson, E. (2006). *Interactions with context*. Ph.D. Dissertation, MIT. http://www-personal.umich.edu/~ericsw/Swanson,%20Interactions%20with%20Context.pdf.
78. Swanson, E. (to appear). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*.
79. Unwin, N. (1999). Quasi-realism, negation and the Frege-Geach problem. *The Philosophical Quarterly*, *49*, 337–352. doi:10.1111/1467-9213.00146.
80. Unwin, N. (2001). Norms and negation: a problem for Gibbard's logic. *The Philosophical Quarterly*, *51*, 60–75. doi:10.1111/1467-9213.00214.

81. van Benthem, J., & Liu, F. (2007). Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, *17*, 157–182.
82. Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, *25*, 221–261. doi:10.1007/BF00248150.
83. von Fintel, K., & Iatridou, S. (2009). *Covert modals? One particular case*. Lecture delivered at LSA. Berkeley. http://web.mit.edu/fintel/lsa220-class-6-handout.pdf.
84. Vranas, P.B.M. (2008). New foundations for imperative logic I: logical connectives, consistency, and quantifiers. *Noûs*, *42*, 529–572. doi:10.1111/j.1468-0068.2008.00684.x.
85. Vranas, P.B.M. (2010). In defense of imperative inference. *Journal of Philosophical Logic*, *39*, 59–71. doi:10.1007/s10992-009-9108-8.
86. Vranas, P.B.M. (2011). New foundations for imperative logic II: pure imperative inference. *Mind*, *120*, 369–446. doi:10.1093/mind/fzr032.
87. Williams, B.A.O. (1963). Imperative inference. *Analysis*, *23*(1), 30–36. doi:10.2307/3326619.
88. Wittgenstein, L. (1958). *Philosophical investigations*. Oxford: Blackwell.
89. Yalcin, S. (2007). Epistemic modals. *Mind*, *116*, 983–1026. doi:10.1093/mind/fzm983.
90. Yalcin, S. (2011). Nonfactualism about epistemic modality. In B. Weatherson, & A. Egan (Eds.), *Epistemic modality*. Oxford: Oxford University Press.
91. Yalcin, S. (2012). A counterexample to modus tollens. *Journal of Philosophical Logic*, *41*, 1001–1024. doi:10.1007/s10992-012-9228-4.
92. Zimmerman, D. (1980). Force and sense. *Mind*, *89*, 214–233. http://www.jstor.org/pss/2253170.