

An Intention-Based Semantics for Imperatives

Daniel W. Harris*

October 10, 2016

Abstract

I argue that imperative clauses semantically encode directive force, and that performing a directive act is a matter of producing an utterance with the aim of getting one's addressee to form an intention to comply. By implementing these ideas in a formal-semantic theory, I show how to make sense of the ways in which imperatives can combine with declaratives in conjunctions, disjunctions, and conditionals. The resulting theory predicts a range of data about imperative inference, including both the badness of Ross's paradox and the goodness of free-choice inferences. It also shows these inferential data to be an indirect, linguistic manifestation of the coherence norms that govern belief and intention. Finally, I explain the illocutionary variability of imperatives by appealing to a theory of indirect speech acts. The resulting view is a self-standing semantic and pragmatic theory of imperative clauses, but also a case study in the attempt to explain the semantic facts by grounding them in facts about the intentions, beliefs, and practical rationality underlying communication. I hope to show that this research program—'Intention-Based Semantics'—is of more than programmatic interest; it can also yield independently motivated explanations of nuanced semantic phenomena.

My aim here is to defend a theory of the semantics and pragmatics of imperative clauses—clauses like (1), each of the two conjuncts of (2), the first disjunct of (3), and the consequent of (4).

*I have received helpful feedback on earlier versions of some of the ideas presented here from Nate Charlow, Daniel Fogal, Magdalena Kaufmann, Ernie Lepore, Sarah Murray, Stephen Neale, Gary Ostertag, Jim Pryor, Will Starr, and Seth Yalcin. In particular, the work of Charlow (2014) and Starr (ms), as well as conversations with both authors, have had a major impact on this essay. Some of the ideas in this paper originated in my dissertation (2014). This draft is not final, and I would appreciate feedback, which can be emailed to me at danielwharris@gmail.com.

- (1) Do the right thing.
- (2) Fly me to the moon and let me play among the stars.
- (3) Make us omelettes or I'll go get us some bagels.
- (4) Help me if you can.

My central claim is that imperatives semantically encode directive illocutionary force, and that a directive act is constituted by the aim of getting one's addressee to form an intention to comply. A literal use of (1), for example, is one on which the speaker performs a speech act whose aim is for the addressee to form an intention to do the right thing.¹ I give a formal-semantic implementation of this idea by identifying the semantic value of an imperative clause with a type of intention that it can be used literally to produce. I represent intentions as properties of cognitive models, which are set-theoretic representations of agents' states of mind.² I'll show, in §2, how to make sense of the semantic contributions that imperatives make to a variety of complex constructions in which they appear, including conjunctions like (2), disjunctions like (3), and conditionals like (4). In §3, I'll use this theory to explain why some inferences involving imperative clauses, such as (5), strike us as compelling, in something like the way that valid arguments do.

- (5) a. Stop trolling celebrities on twitter!
b. If you don't log off twitter, you're never going to stop.
c. So, log off twitter!

Finally, in §4, I'll show how a theory on which imperatives semantically encode generic directive force can draw on a theory of indirect speech acts in order to explain the wide range of illocutionary acts that we regularly perform when uttering imperative sentences, including both particular subspecies of directing (e.g., requesting, commanding) as well as non-directive acts (e.g., wishing, disinterestedly advising, permitting).

What I'm about to lay out is an *intention-based* theory in two senses. First, as I've already indicated, I identify the semantic value of an imperative clause with the type of intention that one intends to provoke when producing a literal and unembedded

¹As I use the terms, an intention is also a plan, but plans may also be complex structures of intentions.

²Both the idea that imperatives' semantic values are types of plans and the idea of representing them as properties of set-theoretic models of mental states are due to Charlow (2014), whose work is an important source of influence on the present account. I'll point out some of the similarities between Charlow's account and mine, as well as various differences, as I go.

utterance of it. But, second, my title should also be understood as signaling allegiance to Intention-Based Semantics (IBS), the research program founded by Grice (1957) and so-named by Schiffer (1982). The central aim of IBS is to give independently motivated psychological explanations of semantic phenomena. Grice's strategy for this comprised two parts. First, show how to reduce facts about the meanings of linguistic expressions to facts about what speakers mean, or are disposed to mean, when using them. Second, show how facts about speakers mean reduce to facts about the intentions with which they speak. Specifically, Griceans hold that to mean something is to produce an utterance with an *m*-intention (a.k.a., 'meaning intention', 'reflexive intention', or 'communicative intention'), which is an intention to have a psychological effect on an addressee, in part by means of their recognition that one intends to do so. In sum, IBS is a theoretical framework that aims to ground the semantic properties of linguistic expressions in the kinds of effects that speakers *m*-intend to have in using them (or would, if they were to use them).

Different kinds of speech acts, on this view, can be individuated in terms of the different kinds of psychological responses that speakers communicatively intend to produce in their addressees. The effect that is characteristic of a certain type of speech act is what Grice (1968; 1969) calls its '*m*-intended response' or '*m*-intended effect'. Although Grice leaves open the possibility of various kinds of *m*-intended effects, he singles out two for particular scrutiny. What Grice calls '*indicative-type utterances*'—and what I'll call '*assertions*'—are aimed at producing beliefs, whereas what Grice calls '*imperative-type utterances*'—and what I'll call '*directives*'—are aimed at producing intentions to act (Grice, 1968, 1969). So although Grice doesn't formulate his view as a formal-semantic theory, the central claim of this paper is nonetheless already nascent in his work.

Intention-Based Semantics has often been motivated in sweeping, metaphysical terms. Schiffer (1982) argues that, along with a physicalist account of the mind, intention-based semantics holds the promise of finding a place for linguistic meaning in the physical order of things. My aims here are much more modest, but also more precise. The present project is intended as a contribution to descriptive semantics as much as to foundational semantics. By drawing several key ideas from IBS, my hope is to offer satisfying and independently motivated psychological explanations of the data that have fueled the study of imperatives in semantics and pragmatics in recent years. An exciting example of this kind of explanation will come in §3, where I will argue that a wide range of data about imperative inference can be explained in terms of facts about the coherence relations that govern the relationships between beliefs and intentions.

I won't try to explain everything there is to explain about imperatives here. I won't try to show how imperatives' meanings arise compositionally from their parts,

for example. Although I will discuss how my view stacks up against some competing views, moreover, I won't have space here for a full discussion of the many competing theories of imperatives that have been proposed in recent years.³ Instead, I will focus on articulating my positive view and displaying its explanatory power.

1 Cognitive Models

The central building block of my semantics will be the notion of a *cognitive model*, which is an idealized formal representation of an agent's beliefs and intentions. Following Hintikka (1962), I'll model an agent's beliefs as a set of possible worlds. Following Gibbard (2003), I'll model an agent's intentions—or, more broadly, their plans—by means of a set of *hyperplans*. A hyperplan is, in Gibbard's words, 'a maximal contingency plan'—a maximally specific plan for what to do in every possible situation. Hyperplans stand to intentions as possible worlds stand to beliefs. An agent whose beliefs were represented by a set containing a single possible world—call them a 'hyperbeliever'—would be completely opinionated about every possible matter of fact. An agent whose plans were represented by a single hyperplan—a 'hyperplanner'—would be completely decided about what to do in any possible occasion of choice.⁴ Just as an agent's beliefs are modeled by the set of worlds that are compatible with everything that they believe, their plans are modeled in terms of the set of hyperplans that are compatible with what they plan. A cognitive model M_A , which represents the beliefs and plans of an agent A , is a set world-hyperplan pairs, $\{\langle w_1, h_1 \rangle, \langle w_2, h_2 \rangle, \dots\}$. From this notion, we can define A 's belief state—the set of worlds compatible with A 's beliefs—as $B_A = \{w : (\exists h)\langle w, h \rangle \in M_A\}$. Likewise, A 's planning state—the set of hyperplans compatible with A 's plans—is $P_A = \{h : (\exists w)\langle w, h \rangle \in M_A\}$.⁵

³A partial list: Charlow (2010, 2014, 2017); Clark-Younger (2012); Condoravdi and Lauer (2012); Han (1998); Kaufmann (2012, 2013, 2016); Kaufmann and Schwager (2009); Murray (2014); Murray and Starr (MS); Parsons (2012, 2013, 2014); Portner (2004, 2007, 2012, 2017); Roberts (2004, 2015, 2017); Starr (ms, 2010, 2014); von Fintel and Iatridou (ms); Vranas (2008, 2010, 2011, 2016); Zanuttini et al. (2012).

⁴I will be working with a somewhat different conception of hyperplans than Gibbard, in two ways. First, Gibbard's hyperplans determine which actions are permitted and which ruled out on each "occasion of choice", and more than one action may be permitted on some occasions. On my view, each hyperplan picks a single course of action for each occasion of choice. Second, Gibbard's hyperplans are not agent relative; I can entertain hyperplans about what I would do if I were Donald Trump. On my view, hyperplans only specify courses of action in situations that the agent could find themselves in, given what they believe.

⁵The reader may be wondering why cognitive models are sets of world-hyperplan pairs, rather than simply an ordered pair of a belief state and a planning state. The answer is that thinking of cognitive models as single sets rather than pairs of sets will make it easier to define clauses for disjunction and

Gibbard treats hyperplans as primitives, but I will adopt a suggestion from Yalcin (2012, 147–148) and take them to be selection functions that map sets of worlds to subsets of themselves.⁶ To implement this idea, it will help to define a class of agent-relative equivalence relations $R = \{R^A, R^B, \dots\}$ on the set of possible worlds such that, for any worlds, w_1, w_2 , and any agent, A , $w_1 R^A w_2$ if and only if w_1 and w_2 don't differ with respect to how A makes any choice of how to act. This equivalence relation partitions the space of worlds into a mutually exclusive and exhaustive space of equivalence classes, $W_c^A = \{c_1, c_2, \dots\}$, which I will call the space of A 's *choice worlds*.⁷

I use these notions to define the set of hyperplans that are *available* to each agent, meaning the hyperplans that can be included in their planning state.

(6) HYPERPLANS AVAILABLE TO AN AGENT

For any agent A , a hyperplan h is available to A if and only if h meets the following conditions:

(i) $h(W) \in W_c^A$

Gloss: Each hyperplan available to an agent maps the entire space of possible worlds to a single one of A 's choice worlds. Intuitively: a hyperplan that is available to A is one that decides every possible choice about how to act that A could ever face.

(ii) $(\forall x \subseteq W) h(x) = h(W) \cap x$

Gloss: Each hyperplan available to an agent maps subsets of the space of worlds to a corresponding subset of one of the agent's choice worlds. Intuitively: for any belief state that an agent could be in, each hyperplan available to the agent specifies a maximally specific course of action that is compatible with what the agent believes.

A cognitive model M_A is available to A just in case every hyperplan in P_A is available to A .

I will represent a belief as the set of worlds in which the belief is true, and I will represent an agent's belief state as the set of worlds in which all of their beliefs are true. B_A is the intersection of A 's beliefs, and A believes p just in case $B_A \subseteq p$. Representing an agent's intentions (or plans) is slightly more complicated. Again, we can represent a single plan as a set of worlds—intuitively, the set of worlds in

the conditional in a way that is neutral about the types of sub-clauses involves.

⁶My implementation of this idea differs from Yalcin's in roughly the same way that my understanding of hyperplans differs from Gibbard's (see footnote 4).

⁷This way of carving up worlds into choosable chunks is inspired by work on deontic modals; see Kolodny & MacFarlane (2010, fn.28), Charlow (2013), and Fusco (2015, §2).

which the plan is enacted. We can thus define a set I_A , which is the set of worlds in which all of A 's plans are enacted. This set of worlds is given as follows: $I_A = \{w : (\exists h \in P_A)(w \in h(B_A))\}$. If p is the proposition that A will perform the action ψ , then A intends to ψ iff $I_A \subseteq p$, and ψ ing is compatible with A 's plans iff $p \cap I_A \neq \emptyset$.

This way of modeling agents' beliefs and intentions is idealized in two importantly different senses. First, any agent whose beliefs and intentions could be accurately represented by a cognitive model would be ideally structurally rational in various ways, by which I mean that their beliefs and intentions would fit together in ways that obey certain coherence norms. For example, beliefs should be coherent: an agent's beliefs should be logically consistent, and agents should believe things that they take to be consequences of the other things they believe. Likewise, agents' intentions should be coherent: an agent shouldn't both intend to do something and intend not to do it. Moreover, intentions should be doxastically constrained: an agent should not intend to do things that are ruled out by their beliefs. Finally, beliefs and intentions should be means-end coherent: an agent who intends to do so-and-so and believes that doing such-and-such is required for doing so-and-so should also intend to do such-and-such.⁸ Cognitive models, as I've set them up, are designed so as to guarantee these norms. Any agent whose mental state conformed to a cognitive model would therefore be optimally rational in each of these respects.⁹ Obviously, this means that no actual human mind is accurately represented by a cognitive model. But this sort of idealization will play a crucial role in helping us to give an independently motivated explanation of some of the ways in which declarative and imperative sentences interact with one another.

A second, rather different sense in which cognitive models are idealized is that they represent beliefs and intentions in ways that are oversimplified and distorted in various respects. Well-known defects of possible-worlds models of belief include the fact that agents are modeled as automatically believing all necessary truths and disbelieving all necessary falsehoods.¹⁰ Since cognitive models represent an agent's intentions as a subset of the worlds that represent their beliefs, this theory likewise predicts that everyone intends all of the necessary truths. Arguably, these are all

⁸ We can model the state of believing that p is required for q as $\lambda M. [(B_M \cap p) \subseteq q]$. This yields a principle stronger than what is usually meant by strict means-end coherence, since it says nothing about the means happening before the end. In order to capture the weaker interpretation, we would need to add a temporal dimension to cognitive models.

⁹ These coherence norms, or slight variations on them, have been widely defended in the philosophical literature on practical rationality. See, for example, Bratman (1987), Holton (2011), and Broome (2013). For an overview, see Kolodny and Brunero (2013).

¹⁰ Believing all logical truths should be distinguished from believing all necessary truths. The former is an idealization in my first sense whereas the latter is an idealization in my second sense. An ideally structurally rational agent would believe all of the logical truths, but if there are necessary a posteriori truths, then an ideally rational agent needn't believe them.

qualities that empirically adequate models of human minds should not have, even if we assume that the minds in question are perfectly rational. I mention them here in order to set them aside. A more naturalistic model would serve my project here even better, although it would presumably also be significantly more complicated.

The statuses of these two kinds of idealization are rather different. The latter are bugs in the models I’ve constructed—distortions that are the price of simplicity. Idealizations of the former kind are features rather than bugs: they are aspects of the model that will help it to play its role in the semantics that is to come.¹¹

2 A Semantics for Declarative and Imperative Clauses

The language-fragment for which I will construct a semantics contains both declarative and imperative clauses. I will assume that clauses of both types factor, at LF, into a mood-marker (‘ \triangleright ’ for declarative, ‘!’ for imperative) and a sentence-radical, ϕ , so that an arbitrary declarative sentence will be represented as $\triangleright\phi$ and an arbitrary imperative sentence will be represented as $!\phi$.¹² To keep things simple, I will assume that sentence radicals have propositions (modeled as sets of worlds) as their semantic values ($(\forall\phi)(\llbracket\phi\rrbracket \subseteq W)$). For convenience, I will say that the semantic value of a sentence’s radical is the sentence’s (propositional, or semantic) content.¹³

Since I’ll be abstracting away from the semantics of sentence radicals except to assume that they express propositions, a semantics for atomic declarative and imperative clauses amounts to a semantics for two operators, \triangleright and $!$, whose semantic values compose with propositions. For example, I’ll take the semantic value of the imperative clause, ‘do your homework!’, in a given context, to be the result of composing the semantic value of the imperative mood-marker ‘!’ with the proposition

¹¹In the terminology of Michael Weisberg, idealizations of my latter sort are *Galilean* idealizations—distortions intentionally introduced into a model in order to make its subject matter computationally tractable, but which will eventually have to be lifted for the sake of predictive accuracy (Weisberg, 2007, 640–1). By contrast, idealizations of my former sort are *minimalist* idealizations, which aim to isolate just those factors in a system that play a special role in explaining how the target phenomenon works (Weisberg, 2007, 642). Unlike Galilean idealizations, Weisberg argues that there is no pressure to eventually lift minimalist idealizations, since their purpose is not broad predictive accuracy but the isolation of distinct factors within a system. As we’ll see, the purpose of the minimalist idealizations that I’ve taken on here is to isolate the nature and underlying source of imperatives’ logical properties.

¹²Others who assume a distinction between mood-marker and sentence radical include Charlow (2014, 2017); Davidson (1979); Grice (1968); Kaufmann (2012); Starr (ms, 2010). On the other hand, Zanuttini et al. (2012) argue for a syntactic analysis of imperatives on which they cannot be factored into mood-marker and sentence radical. Although I assume the former view here, I suspect that the present semantics could be adapted to the alternative syntax.

¹³I argue elsewhere that sentence radicals have rough-grained types of propositions, rather than propositions, as their semantic values (Harris, 2014, MS). I will ignore this complication here, though it is relatively straightforward to adjust my system to accommodate it.

that α does their homework, where α is the addressee in the context.

The syntax of the language with which I'll be working is given by the following definition:

- If Φ and Ψ are both declarative sentences, then $\text{not-}\Phi$, $(\Phi \text{ and } \Psi)$, $(\Phi \text{ or } \Psi)$, and $(\text{If } \Phi, \Psi)$ are also declarative sentences. Nothing else is a declarative sentence.
- If Φ and Ψ are (declarative or imperative) sentences, then $(\Phi \text{ and } \Psi)$ and $(\Phi \text{ or } \Psi)$ are sentences.¹⁴
- If Φ is a declarative sentence and Ψ is a sentence, then $(\text{If } \Phi, \Psi)$ is a sentence.
- Nothing else is a sentence.

I'll give the semantics of atomic sentences by mapping each one to a property of cognitive models. The possession of this property is the type of effect that a speaker *m*-intends to produce in an addressee when they perform a direct and literal speech act using the sentence—in short, it is the sentence's *m*-intended effect. I will also say that a cognitive model M *satisfies* a sentence Φ just in case M possesses the property specified by Φ 's semantic value—i.e., just in case $\llbracket \Phi \rrbracket(M) = 1$.

I'll begin with the base case of the atomic sentences.

(6) ATOMIC DECLARATIVE CLAUSES

$$\llbracket \triangleright \phi \rrbracket = \lambda M. [B_M \subseteq \llbracket \phi \rrbracket]$$

(7) ATOMIC IMPERATIVE CLAUSES

$$\llbracket !\phi \rrbracket = \lambda M. [I_M \subseteq \llbracket \phi \rrbracket]$$

Intuitively, these clauses say that the *m*-intended effect of uttering a declarative sentence is for the addressee to believe the sentence's content, and the *m*-intended effect of uttering an imperative sentence is for the addressee to intend the sentence's content. In performing a literal and direct speech act with the declarative sentence, 'Grice was right', for example, a speaker intends to get their addressee to believe that Grice was right. In performing a literal and direct speech act with the imperative sentence, 'Do your homework', on the other hand, a speaker intends their addressee to form an intention to do their homework.

¹⁴Note that that '!' never scopes under negation. I see no reason to think otherwise.

2.1 Conjunction and Disjunction

Imperatives can be both conjoined and disjoined with other imperatives.

(8) Attack and don't leave anyone standing.

(9) Take the bus or walk.

More importantly, imperatives can be mixed with declaratives under both conjunction and disjunction.

(10) Buy me a drink and I'll pay you back.

(11) Pay up or the whole world will find out about your pog collection.

Many authors have noted that mixed conjunctions and disjunctions sound like indicative conditionals in disguise. For example, someone uttering (10) could seemingly have gotten the same point across with (12), and one might as well utter (13) as (11).

(12) If you buy me a drink, I'll pay you back.

(13) If you don't pay up, the whole world will find out about your pog collection.

Based on examples like these, some have concluded that mixed conjunctions and disjunctions always have conditional meanings. For example, von Fintel and Iatridou (ms) draw on a wide range of nuanced, cross-linguistic data to argue that mixed conjunctions are instances of a broader pattern of conditional conjunctions, which can also take other forms:

(14) One more eye roll and you're out of here, mister.

(15) He brings that dog to the park and the locals will act snobby.

Based on this analysis, together with the robust generalization that conditional conjunctions otherwise don't allow modals in their first conjuncts, von Fintel and Iatridou (ms) argue that these constructions tell against analyses of imperatives that identify them with deontic modals, such as those of Kaufmann (2012) and Han (1998). I find this line of argument quite compelling, at least when it comes to those mixed conjunctions that do read as conditionals.

I have been convinced by Starr (ms), however, that not all mixed conjunctions and disjunctions are disguised conditionals. To my ears, for example, the most prominent reading of (16) is one on which my finding a table is not meant to be understood as conditional on your buying me a drink; rather, I am merely proposing courses of action for both of us to take.

(16) Buy me a drink and I'll find us a table.

One could get the same point across by saying 'I'll find us a table. Buy me a drink,' for example.

Similarly, Starr imagines a scenario in which A and B are at used book sale, and have each found three books, but only have enough cash to buy five.

- (17) A: Put back Waverly or I'll put back Naked Lunch. I don't care which.
(I'll put back Naked Lunch or you put back Waverly. I don't care which.)
B: I'm fine with either too.

According to Starr, "...this disjunction does not have a negative conditional meaning. This is clear from its reversibility and from the fact that these others cannot be followed with either indicator of a free-choice reading: *I don't care which* and *I'm fine with either too*" (Starr, ms, §2.1.1). This seems quite right to me.

Whereas mixed conjunctions with conditional meanings make trouble for theories that posit modals in the LFs of imperatives, mixed conjunctions and disjunctions that don't have conditional meanings make trouble for another class of theories, according to which imperatives have non-modal, non-propositional semantic values. Portner (2004; 2007; 2012; 2017) and Zanuttini et al. (2012) have argued that imperatives' semantic values are properties, for example, and this proposal has been taken up in modified forms by both von Stechow and Iatridou (ms) and Roberts (2015, 2017). But, as Starr (ms) points out, this makes it impossible for them to explain genuine cases of mixed conjunction and disjunction, given the widely held view that only expressions of the same semantic type can be conjoined and disjoined. This problem is not a straightforward one for these accounts to solve, since they explain the distinctive speech-act potentials of clauses of a certain type in terms of the assumption that different types of clauses also have different semantic types (Portner, 2004).

If all of this is right, then we need semantic clauses for 'and' and 'or' that allow true mixed readings, but without identifying imperatives with modals. My strategy, is to hold that the semantic value of a mixed-clause sentence is an property of cognitive models that can be defined as a function of the properties of cognitive models that serve as the semantic values of its clausal parts. Intuitively: the effect one m-intends to produce with a complex, multi-clause sentence is a function of the effects one would m-intend to produce in uttering its parts.

Conjunction is the simplest case. Normally, in uttering the conjunction of two clauses, the speaker intends to produce both of the m-intended effects associated with the two clauses. The m-intended effect of 'dogs bark and cats meow' is a state of believing that dogs bark and believing that cats meow. The m-intended effect of 'Attack and don't leave anyone standing' is a state of intending to attack and intending

not to leave anyone standing. The m-intended effect of (the genuinely conjunctive reading of) ‘Buy me a drink and I’ll find us a table’ is the state of intending to buy the speaker a drink and believing that the speaker will find a table. We can capture these and more complex conjunctions with the following clause:

$$(18) \text{ CONJUNCTION} \\ \llbracket \Phi \text{ and } \Psi \rrbracket = \lambda M . \llbracket \Phi \rrbracket(M) \text{ and } \llbracket \Psi \rrbracket(M)$$

What kind of speech act corresponds to this kind of m-intended effect? The answer is that speech acts of this kind are neither directives nor assertions, but hybrids of the two. Other constructions that mix imperatives with declaratives will be associated with other kinds of directive–assertion hybrids. The fact that there are no colloquial names for such speech acts in English should not deter us from positing them.

Disjunction is a bigger challenge. It might at first be tempting to simply repurpose the clause for conjunction, replacing the metalinguistic ‘and’ with an ‘or’. But it would follow from such an account that satisfying a disjunction requires satisfying at least one of its disjuncts. The m-intended effect of ‘Either Fido is a good dog or Fido will bite’ would be a state of either believing that Fido is a good dog or believing that Fido will bite. But this is not the right result: in believing a disjunction, one needn’t believe either of the disjuncts. Instead, one must be in a state with the following characteristics: (a) every possibility compatible with one’s beliefs entails at least one or the other of the disjuncts, and (b) someone in this state whose beliefs changed in such a way as to rule out one of the disjuncts would, *ceteris paribus*, come to believe the other. Disjunctions involving imperatives have analogous features. If I am successful in changing your mind when I utter ‘Get pilsner or get IPA’, it needn’t be the case that you form an intention to get pilsner or that you form an intention to get IPA; you need only enter a state of mind that is incompatible with intending to get neither.

How must a cognitive model be structured in order to represent minds with these properties? In order to conform to the m-intended effect of a disjunction, a cognitive model M needn’t satisfy either disjunct as a whole, but there must be two subsets of M , each of which satisfies a different disjunct, and the union of these subsets must be M . We can state this formally as follows:¹⁵

$$(19) \text{ DISJUNCTION} \\ \llbracket \Phi \text{ or } \Psi \rrbracket = \lambda M . (\exists M^a)(\exists M^b) \\ \llbracket \Phi \rrbracket(M^a) \text{ and } \llbracket \Psi \rrbracket(M^b) \text{ and } M^a \cup M^b = M]$$

This clause handles mixed imperative–declarative disjunctions as well as disjunctions with uniform disjuncts. In uttering (17)a with its true disjunctive sense, for

¹⁵This clause for disjunction is preliminary; I will defend a slight modification of it in §3.2.

example, I intend to put my addressee into a mental state that is incompatible with both failing to intend to put back Waverly and failing to believe that I will put back Naked Lunch.

(17) a. Put back Waverly or I'll put back Naked Lunch.

A plausible purpose of putting someone into this sort of mental state is to facilitate further joint planning. This state of mind is such that, in order to form a rational intention to keep Waverly, my addressee must satisfy themselves that I will put back Naked Lunch; forming the latter belief in a rational way, in turn, requires action on their part. On the other hand, upon forming the mental state that is the semantic value of (17), if my addressee subsequently forms the belief that I am not going to put back Naked Lunch, then they are rationally committed, *ceteris paribus*, to form an intention to put back Waverly. This seems the right way to think about what someone would normally be trying to accomplish when uttering (17)a.

2.2 Conditionals

The plot further thickens when we try to account for conditional imperatives (CIs).

(20) Attack if the weather is good.

(21) If the bartender comes back, buy me a drink.

In uttering a CI, the speaker issues what might be called a hypothetical directive—a directive that must be acted on only once some condition is met. In order to comply with (21), for example, one shouldn't attempt to buy the speaker a drink immediately, but merely once one forms the belief that the bartender has come back.

A crucial fact, for our purposes, is that the consequents of conditionals may be mixed conjunctions and disjunctions, such as those discussed above.

(22) If the bartender is back, buy me a drink and I'll find us a table.

(23) If we only have enough money for five books, put back Waverly or I'll put back Naked Lunch. (I don't care which.)

This is a particularly striking datapoint, as it rules out the possibility of giving distinct semantic clauses for conditional imperatives and indicative-conditionals, which is the most commonly adopted approach to CIs (Charlow, 2014; Roberts, 2015; von Stechow, ms). What we need is a single conditional that can handle declarative, imperative, or mixed consequents.

The guiding idea of my account is that the aim of a conditional speech act is to put one’s addressee into a state of mind that will rationally compel them, *ceterus paribus* to satisfy the consequent should they come to satisfy the antecedent. For example, in uttering ‘If Fido is lost, we’re in trouble’, my aim is to put my addressee into a state such that, if they are in this state and come to form the belief that Fido is lost, they will be rationally compelled to form the belief that we’re in trouble. And, in uttering the conditional imperative, ‘If the bartender comes back, buy me a drink’, I intend to put my addressee into a state such that, if they subsequently form the belief that the bartender has come back, they will be rationally compelled to form an intention to buy me a drink.

We can formulate a clause for a conditional of this kind more easily if we first define the notion of a maximal submodel:

(24) MAXIMAL SUBMODEL

A maximal Φ -satisfying submodel M^Φ of a model M is a model that meets all of the following conditions:

- (i) $M^\Phi \subseteq M$,
- (ii) $\llbracket \Phi \rrbracket(M^\Phi)$, and
- (iii) there is no model M^* such that $M^\Phi \subset M^* \subseteq M$ and $\llbracket \Phi \rrbracket(M^*)$.

With this notion in hand, we can say that a model satisfies a conditional if and only if its maximal antecedent-satisfying submodels satisfy the consequent.¹⁶ Formally:

(25) CONDITIONAL

$\llbracket \text{If } \Phi, \Psi \rrbracket = \lambda M . (\forall M^\Phi)(\llbracket \Psi \rrbracket(M^\Phi))$

The intuitive idea here is that the m-intended effect of uttering a conditional is a state of mind that, were it to change in a minimal way (or possibly not at all) so as to satisfy the antecedent, would also satisfy the consequent.¹⁷

It will help to work through a simple example to see how this clause handles conditional imperatives. Suppose that you and I are about to go into the bar where your friend, Quinn, sometimes works. When he’s working, he gives you special deals on drinks. We don’t know whether Quinn will be working tonight. So, you’re

¹⁶This way of understanding conditionals is influenced by Kolodny & MacFarlane (2010), who are, in turn, influenced by Cantwell (2008). There are several differences between my conditional and theirs, the most notable being that their conditionals cannot have imperatives in their consequents. Another close predecessor of the present account of imperative conditionals is Charlow’s (2014), but Charlow gives distinct semantic clauses for declarative conditionals and imperative conditionals, and so his account cannot accommodate conditionals with mixed imperative–declarative consequents.

¹⁷It is worth pointing out that this proposal does not extend to counterfactual conditionals

currently in a state of ignorance about whether Quinn will be there, and your plans are currently compatible with either buying the first round or not buying the first round, whether or not Quinn is working. Abstracting away from the rest of your beliefs and plans, we can represent your state of mind with the cognitive model M^- . (Here, QR is the proposition that Quinn is working and you buy the first round, Qr is the proposition that Quinn is working and you don't buy the first round, and so on.)

- (26)
- $M^- = \{\langle QR, h_1 \rangle, \langle Qr, h_1 \rangle, \langle qR, h_1 \rangle, \langle qr, h_1 \rangle, \langle QR, h_2 \rangle, \langle Qr, h_2 \rangle, \langle qR, h_2 \rangle, \langle qr, h_2 \rangle, \langle QR, h_3 \rangle, \langle Qr, h_3 \rangle, \langle qR, h_3 \rangle, \langle qr, h_3 \rangle, \langle QR, h_4 \rangle, \langle Qr, h_4 \rangle, \langle qR, h_4 \rangle, \langle qr, h_4 \rangle\}$
 - $B^- = \{QR, Qr, qR, qr\}$
 - $P^- = \{h_1, h_2, h_3, h_4\}$
 - $h_1(B^-) = \{QR\}$
 - $h_2(B^-) = \{Qr\}$
 - $h_3(B^-) = \{qR\}$
 - $h_4(B^-) = \{qr\}$
 - $I^- = \{QR, Qr, qR, qr\}$

Now, suppose that I utter (27) while we're still outside the bar.

- (27) If Quinn is working, buy the first round.

According to the present semantics, my aim in producing this utterance is to put you into a mental state that could be modeled by a cognitive model M that satisfies the following property: any maximal submodel of M that satisfies 'Quinn is working tonight' also satisfies 'buy the first round'. In fact, there are multiple ways to change M^- to give it this property—a point to which I'll return momentarily. But one simple way to change M^- so that it satisfies (27) would be to kick out all of the points containing h_2 , thereby turning it into M^+ :

- (28)
- $M^+ = \{\langle QR, h_1 \rangle, \langle Qr, h_1 \rangle, \langle qR, h_1 \rangle, \langle qr, h_1 \rangle, \langle QR, h_3 \rangle, \langle Qr, h_3 \rangle, \langle qR, h_3 \rangle, \langle qr, h_3 \rangle, \langle QR, h_4 \rangle, \langle Qr, h_4 \rangle, \langle qR, h_4 \rangle, \langle qr, h_4 \rangle\}$
 - $B^+ = \{QR, Qr, qR, qr\}$
 - $P^+ = \{h_1, h_3, h_4\}$
 - $h_1(B^+) = \{QR\}$

- $h_3(B^+) = \{qR\}$
- $h_4(B^+) = \{qr\}$
- $I^+ = \{QR, qR, qr\}$

In this state of mind, you're still agnostic about whether Quinn will be working, and about whether you'll buy the first round. Moreover, it's still compatible with what you believe that Quinn will be working and you won't buy the first round. (We can suppose that you're uncertain whether you have enough cash.) All that's changed is that now you have a new contingency plan, a hypothetical intention to buy the first round if Quinn is working.¹⁸

The above example illustrates a feature of this semantics that sets it apart from dynamic-semantic theories in which sentences' semantic values are context-change potentials—functions that map each possible state of the context to a uniquely updated, post-utterance state. In Starr's (ms) system, for example, the semantics spells out a unique way in which the context must be updated with a new preference when an imperative is uttered. In the present system, by contrast, there may be more than one way in which an addressee's mind can change in order to comply with a speech act. In the case just discussed, for example, you could alternatively have satisfied (27) by forming the belief that you would buy the first round, in addition to a plan to do so. Whether you update in this way, or merely with an unconfident plan should depend on factors that are extraneous to semantics, such as your degree of confidence that you can afford to buy us a round. So, while this semantics is built around the idea that speech acts are aimed at changing the minds of their addressees, it admits of some flexibility in how the addressee does so.

This point is closely related to one made by Charlow (2014, §5.6.2), who argues that this sort of flexibility constitutes a significant advantage over dynamic-semantic theories. Charlow's version of the point centers around the fact that we are sometimes on the receiving end of assertions that conflict with our current beliefs, and directives that conflict with our current plans. But updating a state of mind with conflicting input is a complex, inferential, and non-monotonic process that is part of the subject matter of epistemology and artificial intelligence, not formal semantics. Proponents of both dynamic semantics and dynamic pragmatics have avoided this issue mainly by idealizing away from incoherent conversational contexts. For example, as Stalnaker is the first to admit, his theory of context stops making predictions about what happens in "defective contexts", when a group of interlocutors

¹⁸ Although some philosophers of action have argued that intending to ψ requires believing that one will ψ (e.g. Grice 1971), most now agree that this requirement is too strong. For example, I am writing this paper with the intention of publishing it in journal X, but it is compatible with my beliefs that I will wind up publishing it in journal Y instead, or that I will contribute it to an edited volume. Cognitive models are designed to capture this looseness in the constraints that beliefs place on intentions.

take inconsistent propositions for granted—at least when this is relevant to the conversation going forward (1978). Similarly, most dynamic-semantic and dynamic-pragmatic theories of imperatives make either weird predictions or no predictions about what happens when contradictory directives are issued.¹⁹ This idealization is what allows dynamic theorists to think of speech acts in terms of tidy and unique updates. But this idealization away from defective contexts will eventually have to be lifted, because agents sometimes do receive directives that are contrary to their current plans, preferences, or commitments, as well as assertions that conflict with their current information. I therefore agree with Charlow that semantics should be concerned only with thin descriptions of the effect at which we aim when uttering a sentence, allowing for flexibility in how this effect is integrated into addressees' minds.

3 Inference and Infelicity

Some inferences involving imperatives strike us as compelling in something like the way that valid inferences do.

- (29) a. Take out the trash and mow the lawn!
 b. So, take out the trash! (Charlow, 2014, §2.2)
- (30) a. Attack at dawn if the weather is fine.
 b. The weather is fine.
 c. So, attack at dawn! (Parsons, 2013, 61)
- (31) (a) Buy me an IPA!
 (b) If you don't go to the bar, you won't buy me an IPA.
 (c) So, go to the bar!

Semantic theories are usually held responsible for explaining our intuitions about valid-seeming patterns of inference like these, and entailment judgments often provide crucial empirical footholds for semantic theories. But it is clear that our usual strategies for offering this sort of explanation will have to be revised in some way. One way to appreciate the puzzle posed by these arguments is to consider the following inconsistent triad (Clark-Younger, 2012; Parsons, 2014):

- (32) (i) There are non-trivially valid arguments containing imperatives.
 (ii) Imperatives are not truth-apt.
 (iii) Validity is truth-preservation.

¹⁹For example, Portner's theory predicts that contrary directives cancel each other out (2012); but von Stechow and Iatridou (ms) point out that this is not the right prediction.

Each of (32)i–iii has some support. The support for (32)i comes from the intuition of validity that we get from some arguments involving imperatives, such as the above examples, which seems a lot like the intuition of validity that we get from regular declarative arguments. The non-truth-aptness of imperatives is not only highly intuitive, but also motivated by a wide range of considerations that I won't repeat here (Charlow, 2014). Meanwhile, the support for (32)iii comes, at least in part, from the weight of a tradition extending back to Bolzano, Frege, and Tarski, according to which logic is the study of truth, and the consequence relation is the study of truth preservation.

Two particularly difficult data for a theory of imperative inference to explain are ROSS'S PARADOX, which raises the problem of explaining why (33)c doesn't seem to follow from (33)p, and FREE CHOICE, which raises the problem of explaining why (34)ci and (34)cii *do* seem to follow from (34)p (though only when (34)ci and (34)cii are being used to acquiesce or permit, and not when they are understood as directives).

- (33) ROSS'S PARADOX
 (p) Post the letter.
 (c) $\not\models$ Post the letter or burn the letter.
- (34) FREE CHOICE
 (p) Have some tea or have some coffee.
 (ci) \models Have some tea.
 (cii) \models Have some coffee.

Since neither of these inference patterns is classically valid, it seems that our theory of imperative inference will have to be non-classical.

Finally, in addition to apparent instances of the consequence relation that involve mixtures of imperatives and declaratives, there are also apparent cases of mixed imperative–declarative inconsistency (Starr, ms, §2.2):

- (35) # Unicorns have never existed, and never will. Bring me a unicorn!
- (36) # The door is open. Open the door!

More generally, if it is clear that the addressee can't ψ , for whatever reason, and p is the proposition that the addressee ψ s, then it will sound inconsistent or infelicitous to utter the imperative $!p$. An adequate account of the semantics and pragmatics of imperatives must explain this sort of information-dependence—something that is difficult for many extant accounts to do, as Starr (ms) points out.

3.1 A Theory of Imperative Inference

With the help of the semantics laid out in §2, I can explain much of this data by defining a relation \vdash as follows.

- (37) INFERENCE
 $\{\Phi_1, \dots, \Phi_n\} \vdash \Psi$ iff:
 $(\forall M)$ if $\llbracket \Phi_1 \rrbracket(M), \dots, \llbracket \Phi_n \rrbracket(M)$, then $\llbracket \Psi \rrbracket(M)$

In English: a conclusion can be inferred from a collection of premises just in case any cognitive model that satisfies all of the premises also satisfies the conclusion. The intuitive idea here is that an argument will appear valid to us when its conclusion would produce no effect (or a redundant effect) in a structurally rational mind that already satisfies its premises.

Several well known argument forms, including and-elimination, modus ponens, and disjunctive syllogism, are instances of \vdash , irrespective of whether the sentences involved are declarative, imperative, or combinations of the two. This allows us to predict a wide range of intuitive inferential data, including the fact that (29) and (30) strike us as compelling.

What's more interesting than the predictions themselves, I think, is the fact that the account laid out so far makes it possible to say something about *why* \vdash works the way it does. To see what I mean, it will be helpful to consider a pair of possible answers to the question of how the present account resolves the inconsistent triad (32). The first option would be to think of \vdash as a redefinition of logical consequence. On this view, consequence can no longer be thought of as a relation of truth-preservation, since it is defined over sentences that aren't truth-bearers. This view resolves the triad by rejecting (32)iii.

One reason not to think of \vdash as the relation of logical consequence is the desire to preserve the traditional idea that real logical consequence is about truth preservation. A related reason is that this proposal would commit us to psychologism about logical consequence, according to which logic is the study of some facts about human psychology. After all: \vdash is defined in terms of cognitive models, which are idealized representations of human minds. As a proponent of Intention-Based Semantics, I think it is a good idea to psychologize the semantic properties of linguistic expressions. But psychologism about logic does not follow, and I think there are good reasons to be skeptical about that position.²⁰

So, on to the second way of understanding of how \vdash resolves the triad. I prefer to understand \vdash as a model of ways in which the coherent norms governing our

²⁰For some classic attacks on psychologism about logic, see Frege (1894); Husserl (1975). For a more recent survey, see Pelletier et al. (2008).

mental states are reflected in the language we use to provoke those mental states. The reason that certain arguments involving natural-language sentences strike us as compelling, I submit, is that any structurally rational mind that exemplifies the semantic values of the premises would already exemplify the semantic value of the conclusion. Our intuitions about the validity of natural language arguments are, on this view, an indirect manifestation of our sensitivity to facts about what it takes to be a structurally rational agent.

The view I'm now advocating resolves the triad by rejecting (32)i. We can continue to think of logical consequence as truth preservation, but should understand truth preservation as a relation on propositional contents, first and foremost, and only secondarily on the sentences whose function is to provoke beliefs with those contents. Since one of the norms governing the coherence of mental states is that we should believe the logical consequences of our beliefs, \vdash will be closely related to logical consequence when we restrict our attention to inferences involving only declarative clauses.

But logical closure is not the only norm governing the rationality of our mental states. There are also coherence relations governing non-doxastic mental states, and some that govern the connections between beliefs and mental states of other kinds, including intentions. In §1, I pointed out that cognitive models are designed to guarantee several of these independently motivated coherence relations. We are now in a position to recognize the explanatory payoffs of setting things up in this way. Namely: each of the coherence-guaranteeing features of cognitive models that I outlined earlier plays a role in predicting and explaining data about imperative inference and infelicity.

Consider, for example, (31), which turns out to be an instance of \vdash .

- (31) (a) Buy me an IPA!
- (b) If you don't go to the bar, you won't buy me an IPA.
- (c) So, go to the bar!

The intuition that this inference is valid is a reflex, bubbling up through our semantic competence, of our sensitivity to the norm of strict means-end coherence. Any agent who intends to buy me an IPA and believes that buying me an IPA requires going to the bar is rationally required, by dint of this norm, to intend to go to the bar. The present semantics captures this and other instances of the same form of inference by guaranteeing the strict means-end coherence of cognitive models. It thereby generates not just an accurate prediction about inferences like (31); it also explains why we find such inferences compelling by appealing to an independently motivated principle governing our psychological states.

Cognitive models are also designed to guarantee that agents' intentions are constrained by their beliefs. In the present model, this amounts to the fact that if an

agent's belief state rules out possibilities in which they will ψ , it follows automatically that they do not intend to ψ . (Formally: for any M , A , and p , if $B^A \cap p = \emptyset$, then $I^A \not\sqsubseteq p$.) This property of cognitive models allows us to predict the infelicity of (35) and (36).

(35) Unicorns have never existed, and never will. # Bring me a unicorn!

(36) The door is open. # Open the door!

The two halves of (35) cannot be satisfied by the same cognitive model. On the assumption that it is generally infelicitous to perform two speech acts with incoherent aims in close succession, the infelicity of (35) is explained. The same goes for (36), given the background assumption that open doors can't be opened.

In the real world, of course, intuitions like these, about both the felicity of speech acts and the validity of inferences, will be only as robust as our sensitivity to the coherence norms at issue. In complex cases, these norms may go unnoticed altogether, and this explains why we lack intuitions about complex arguments. This can be explained by appeal to a kind of competence–performance distinction: although we are generally sensitive to psychological coherence norms, we may lack the psychological resources to appreciate all of the implications of these norms, particularly in complex cases. Cognitive models idealize away from performance restrictions of this kind.

3.2 Ross's Paradox and Free Choice

Finally, what about the inference patterns that cause the most trouble for logics of imperatives, since they seem to require departures from classical inference rules? First, consider ROSS'S PARADOX.

- (33) ROSS'S PARADOX
 (p) Post the letter.
 (c) \neq Post the letter or burn the letter.

The system I have constructed fails to make the right prediction here, as it validates inferences of the form $!\phi \vdash !\phi$ or $!\psi$. However, a small alteration to the semantics for disjunction given in §2 would allow us to block this inference. The altered clause reads as follows.

- (38) STRONG DISJUNCTION

$$\llbracket \Phi \text{ or } \Psi \rrbracket = \lambda M. [(\exists M^a : I_{M^a} \neq \emptyset)(\exists M^b : I_{M^b} \neq \emptyset)$$

$$\llbracket \Phi \rrbracket(M^a) \text{ and } \llbracket \Psi \rrbracket(M^b) \text{ and } M^a \cup M^b = M]$$

All that's been added here are the restrictions on the existential quantifiers. In effect, these restrictions guarantee that the m-intended effect of a disjunction must be non-trivially disjunctive: the speaker intends the hearer to enter a state of mind that is compatible with both disjuncts. This strengthening blocks ROSS'S PARADOX because even if you have an intention to post the letter, it does not follow that both posting the letter and burning the letter are compatible with your intentions.

How plausible is STRONG DISJUNCTION as a semantics for English clausal 'or'? It may be tempting to balk at this strengthening, and to prefer the weaker version presented in §2. It is intuitive that an addressee can satisfy a disjunctive imperative by forming a plan to comply with just one of the disjuncts. If I address an utterance of (39) to you and you form an intention to wash the car, then it seems that you've complied with my request.

(39) Mow the lawn or wash the car.

An analogous point is tempting when applied to disjunctions of declaratives. Suppose I address an utterance of (40) to you.

(40) Either Santa Claus is real or my parents have been lying to me.

It seems that one could accept what the speaker has said merely by believing that their parents have been lying to them, in this case. If these intuitions are to be trusted, then the original, weak semantics for disjunction should be preferred.

Let me try to make the case for STRONG DISJUNCTION. The motivating idea of this defense is that, in uttering a disjunction, the speaker intends for their addressee to treat both disjuncts as live options, even if they do so only temporarily before settling on one or the other. In cases where a hearer quickly chooses one or the other, they won't have done as the speaker intended unless they reached this state by taking both options seriously and deliberating about which to settle on.

To see this, consider some of the different roles that uttering a disjunction of imperatives, such as (39), could play in a discourse. Perhaps the easiest-to-imagine scenario is one in which the addressee has neither an intention to mow the lawn nor an intention to wash the car, and the speaker utters (39) in order to get them to choose at least one of these two courses of action and form an intention to do it. But, for this choice to be genuine, the addressee must first entertain both plans for the purposes of engaging in practical reasoning. On this way of thinking about things, if you comply with my utterance of (39) by forming an intention to wash the car, this should be a downstream consequence of initially making both mowing the lawn and washing the car and mowing the lawn compatible with your plans, and then choosing between them.

Now consider a different scenario, in which B utters (39) as part of the following dialogue.

- (41) A: I'm going to wash the car.
B: Wash the car *or* mow the lawn.

Here, A makes it clear with their utterance that they intend to wash the car. So, if disjunction were weak, then B's reply would be redundant: A's intention to mow the lawn would already satisfy B's whole disjunction by satisfying one of its disjuncts. But B's utterance doesn't seem redundant. Intuitively, B's aim is to for A to at least consider mowing the lawn as an alternative to washing the car. In a scenario in which B has some sort of power over A's actions, B's utterance sounds like an offer of permission to mow the lawn instead of washing the car. Alternatively, B might be offering A advice, or modifying an earlier request. Any of these speech acts can be explained if we adopt STRONG DISJUNCTION: B's utterance can function as a way to get A to reconsider their plans if satisfying a disjunction requires taking both disjuncts to be live options, even if only momentarily and for the purposes of practical reasoning.

A similar conversational move can be made with a declarative disjunction.

- (42) A: The talk will be in room 507.
B: The talk will be in room 507 *or* it will be in room 721.

Here, B's utterance is not redundant, but seems aimed at expanding A's belief state to make it compatible with the second disjunct as well as the first. Again, we can easily explain this reading if we adopt STRONG DISJUNCTION, which entails that A can't satisfy B's disjunction merely by believing that the talk will be in room 507. Although these data shouldn't be taken as decisive—we might try to explain both them and ROSS'S PARADOX pragmatically, for example—they give us some prima facie reason to strengthen our account of disjunction in the way proposed.

Next, consider FREE CHOICE.

- (34) FREE CHOICE
(p) Have some tea or have some coffee.
(ci) \models Have some tea.
(cii) \models Have some coffee.

Before we can settle on a final account of this phenomenon, we need a clearer understanding of weak uses of imperatives—speech acts that are sometimes called 'offers', 'permissions', 'acquiescences', or 'suggestions'. As already noted, instances of FREE CHOICE strike us as compelling only when the conclusions (and, optionally, also the

premises) are understood as weak speech acts rather than strong (directive) acts. This is the most natural reading of (34), whereas other instances of the same pattern don't look as good when their conclusions are most naturally understood as, say, commands. For example:

- (43) (p) Pay the toll or turn around!
 (c) \neq So, turn around!

So, our treatment of FREE CHOICE hangs on our treatment of weak uses of imperatives. And since it is controversial whether weak readings should be given a semantic explanation or a pragmatic one, it is correspondingly unclear whether FREE CHOICE is a semantic or a pragmatic phenomenon.

I have characterized strong uses of imperatives as directives—acts aimed at getting one's addressee to form an intention. Weak uses of imperatives have a different sort of aim. Some examples:

- (44) Have a cookie. [offer]
 (45) (Fine.) Borrow my car. [acquiescence]

In stereotypical utterances of these sentences, the speaker's aim is not to instill a plan of action in the addressee. Rather, the speaker is attempting merely making a new plan of action available, which the addressee is then free to choose. I may achieve my aim when offering someone a cookie despite their never forming a plan to have a cookie, for example, since my aim is merely that they choose whether or not to have a cookie of their own volition. The question, then, is how best to explain the possibility of using an imperative to achieve this sort of effect.

One simple way to explain weak uses of imperatives is to posit a semantic ambiguity and formulate a second semantic clause for weak uses, which I'll follow various others in annotating ' $i\phi$ ' (Charlow, 2014; Starr, ms). The intuitive contrast between strong and weak uses is this: whereas $!\phi$ is used to change the addressee's plans in such a way that they intend $\llbracket\phi\rrbracket$, $i\phi$ is used to change the addressee's plans in such a way that they become compatible with $\llbracket\phi\rrbracket$ (cf. Charlow 2014).

- (46) ATOMIC WEAK-IMPERATIVE CLAUSES
 $\llbracket i\phi \rrbracket = \lambda M. [I_M \cap \llbracket \phi \rrbracket \neq \emptyset]$

Provided that we also adopt STRONG DISJUNCTION as our semantic clause for 'or', it follows that the following inference forms are instances of \vdash :

- (47) $(!\phi \text{ or } !\psi) \vdash i\phi$

(48) $(i\phi \text{ or } i\psi) \vdash i\phi$

In other words, FREE CHOICE is validated.

There are a couple of loose ends to consider. First, it may be desirable to offer a pragmatic explanation of weak uses of imperatives. On this view, these weak uses would be speech acts whose m-intended effects can be given as in (46), but this non-directive force would not get its own semantic clause. I will briefly consider how weak readings could be explained pragmatically in §4.

A second, and more troubling issue, is the question of whether I have misrepresented the LFs of the disjunctions involved in both ROSS'S PARADOX and FREE CHOICE. I have assumed that 'or' takes wide scope over the '!' in both of these inference patterns. But mightn't it be better to represent their LFs as $!\phi \vdash !(\phi \vee \psi)$ and $!(\phi \vee \psi) \vdash i\phi$, respectively? I am not sure. If so, then the semantics presented here would have to be supplemented with, at least, a revised theory of how 'or' works within sentence radicals that scope under '!'. Doing that would go beyond the scope of the present essay, which has been focused only on semantic phenomena that occur at the level of the clause.

Even if the account I've offered here needs to be further developed in this way, however, I think there is reason to be optimistic about the broader strategy of explaining the inferential data in terms of the coherence relations governing mental states. I am led to this conclusion by the fact that both the FREE CHOICE and ROSS'S PARADOX data, like the other inference data I've considered, appear to be grounded in facts about the rational structure of practical reasoning. Assuming that I am rational, if I intend to have tea or coffee, it follows that it is compatible with my plans to have tea. This is the psychological fact that ultimately explains why the FREE CHOICE inference strikes us as valid. Likewise, a rational agent who intends to post the letter doesn't, *ipso facto*, intend to post or burn the letter. This is the psychological fact underlying the felt invalidity of ROSS'S PARADOX. These parallels between the inferential roles of imperatives and the coherence requirements on intentions strike me as deep and important, and lend significant support to the idea that the meanings of imperative clauses should be understood in terms of the intentions that we use them to provoke.

4 Imperatives and Indirect Speech Acts

A final explanandum that I wish to address is the fact that imperatives are regularly used to perform a diverse range of speech acts, including both species within the directive genus and a variety of non-directive acts.

Commands are one species of directive. Intuitive examples of commands include the act performed by a manager who addresses (49) to a subordinate, or some-

one uttering (50) while aiming a gun at their addressee.

(49) Get me that report.

(50) Put your hands up.

Requests form another subcategory of directives. Addressed to a friend who is helping with a renovation, (51) would normally be understood as a request rather than a command, for example.

(51) Pass me the hammer.

The distinction between commands and requests is pre-theoretically intuitive, and the two kinds of acts have different motivational and normative profiles. Requests can normally be refused much more easily than commands, and speech acts of the two kinds present addressees with different kinds of reasons for acting. Moreover, acts of the two kinds exhibit some grammatical differences. Adding 'please' to an imperative forces a request reading, whereas an imperative will sound like a command when followed by an 'or-else'-clause that makes a conditional threat explicit.

(52) Pass me the hammer please.

(53) Put your hands up, or I'll charge you for resisting arrest.

Nothing that I have said in the previous sections allows us to distinguish between requests, commands, and other directives. On my view, imperatives semantically encode what might be thought of as generic directive force. Both requests and commands fall under this rubric, since they're both aimed at getting the addressee to form an intention to comply. But it would be nice to know what distinguishes different species within the directive genus.

There are also several species of non-directive speech act that we typically perform with imperatives. In addressing (54) to a friend in the hospital, for example, I can't rationally intend to change their plans, since I know that they lack control over the outcome.

(54) Get well soon!

(Condoravdi and Lauer, 2012, 39)

Instead, it looks like this use of (54) is merely a way of expressing a wish or preference for my friend to improve. Similarly, there are uses of imperatives that are sometimes referred to as 'disinterested advice' uses, as in B's half of (55).

- (55) A: Excuse me, how do I get to San Francisco?
B: Take the train that leaves from over there in 10 minutes. [points to train station] (Condoravdi and Lauer, 2012, 40)

Some uses of imperatives that we would intuitively classify as advice also seem to be directives. For example, if a friend addresses (56) to me in an urgent tone, it seems clear that their goal is to sway my action.

- (56) Don't quit your job to become a mime!

However, it is consistent with an exchange like (55) that B don't care one way or another about whether A complies. For all B knows, A is just wondering about where San Francisco is, and isn't even trying to get there. In this case, B seems to be uttering the imperative in order to impart information to A—an observation that is supported by the fact that B's utterance constitutes an answer to A's question.

A further class of non-directive uses of imperatives comprise the weak uses of imperatives discussed in §3.2, including permission or invitation uses, as in 1, as well as what (von Fintel and Iatridou, ms, 5) call 'acquiescence' uses, as in (58), and 'indifference' uses, as in (59).

- (57) Have a cookie.

- (58) A: It's getting warm. Can I open the window?
B: Sure. Go ahead. Open it!

- (59) Go left! Go right! I don't care.

What makes these speech acts weaker than directives is that a speaker who performs one can seemingly achieve their aims even if the hearer doesn't form an intention to comply. In §3.2, I explained this fact by positing a second reading of imperatives, on which they are used with the aim of making their content compatible with the addressee's intentions, and not with a directive aim. But there are some good reasons not to explain this second reading by positing an ambiguity. For example, von Fintel and Iatridou (ms, 5) point out that weak uses of imperatives occur across many languages which are otherwise quite different—a fact that doesn't sit well with an ambiguity theory. It is therefore tempting to look for a more systematic explanation of weak readings.

Condoravdi and Lauer (2012) have argued that non-directive uses of imperatives give us reasons to avoid building directive force into the semantics of imperative clauses. More broadly, the illocutionary diversity of imperatives threatens the systematic semantic link between imperatives and directive force that I have posited.

In defense of this link, I will spend the rest of this section arguing that we can explain this diversity pragmatically, by appealing to a theory of indirect speech acts. In performing a directive act, a speaker normally also indirectly communicates a reason for the addressee to comply. Different species within the directive genus are backed by different kinds of reasons for complying. Non-directive uses of imperatives, on the other hand, are indirect speech acts that one can perform by *making as if* to perform a directive act. I will defend these ideas in §§4.3. First, I will say more about what I mean by indirect speech acts.

4.1 Indirect Speech Acts

Indirect speech acts, on the Grice-inspired view that I prefer, subsume the category of conversational implicatures. A speaker, *S*, conversationally implicates *q*, according to Grice, if *S* attempts to communicate *q* without saying *q* explicitly, with the expectation that their hearer will infer that *S* meant *q* partly on the basis of having recognized what *S* said (or made as if to say), and partly on the basis of the assumption that the speaker is being cooperative.

For present purposes, we should take care to distinguish two kinds of conversational implicature. First, speakers may conversationally implicate one proposition, *q*, by saying another proposition, *p*. In this case, the speaker intends for the hearer to believe both *p* and *q*, but intends to communicate *q* by way of communicating *p*. This sort of implicature is exemplified by Grice's example of saying that there is a garage around the corner in order to implicate that the garage is open and is selling gas (1989, 32). Here, the speaker means both that there is a garage around the corner *and* that it is open and selling gas, but intends the hearer to comprehend the latter indirectly, in part by comprehending the former and in part by assuming that the speaker is being cooperative.

Second, speakers may conversationally implicate *q*, not by saying *p*, but by merely *making as if to say p*. In cases of this kind, the speaker intends their hearer to believe *q* but not *p*. Roughly speaking, a speaker, *S*, makes as if to say *p* just in case they use a sentence that would normally be used to say *p*, with the intention that their addressee, who is working under the assumption that *S* is being cooperative, will consider and reject the assumption that *S* means *p*, and conclude that *S* means something else instead. Grice's example of someone ironically uttering 'X is a fine friend' in order to implicate that X is duplicitous is a clear example of this sort of implicature.

In formulating his theory of implicature, Grice focused exclusively on assertoric speech acts. As Grice uses 'say' and 'implicate', to say something is to perform a direct assertoric speech act, and to implicate something is to perform an indirect assertoric speech act. The m-intended effect of saying *p* is for one's addressee to be-

lieve p , and the m -intended effect of implicating q is for one's addressee to believe q . Grice not only focused on assertoric examples; he also formulated his maxims to cover only assertoric cases. The maxim of quality tells us to avoid speaking falsely and without evidence, for example, whereas the maxim of quantity tells us to be appropriately informative (1989, 26–7). These maxims are category mistakes as applied to questions and directives. Grice recognized this shortcoming. He thought that his theory could be generalized to cover indirect speech acts of other kinds, but he never attempted to show how.

I have stated my maxims as if [the purpose of talk exchange] were a maximally effective exchange of information; this specification is, of course, too narrow, and the scheme needs to be generalized to allow for such purposes as influencing or directing the actions of others. (Grice, 1989, 28)

I won't attempt to formulate a fully general theory of indirect speech acts here, but I will assume that such a theory can be given, and that the general contours would follow Grice's theory of implicature.

In particular, it should be clear that indirect speech acts in general can be divided into those that one performs by *performing* a direct speech act, on one hand, and those that one performs by merely *making as if to perform* a direct speech act, on the other hand. A well worn example of the former kind is indirectly requesting the salt by uttering 'Can you pass the salt?'. As noted by Bach & Harnish (1979, §9.2.1), the fact that one is still asking the question, in addition to making a request, can be seen from the fact that we often respond by saying 'yes' or 'no'. A similar, if less conventionalized example, is indirectly requesting that a stranger in a crowded elevator get off my foot by saying, 'you're standing on my foot'—a sentence that could be used to implicate various other things in other circumstances. On the other hand, when a bodybuilder implicates that their addressee is a weakling by saying, 'do you even lift?', they probably shouldn't be understood as asking a genuine question; rather, they're merely making as if to ask one. The same would apply to a case in which my boss indirectly ordered me to work late by saying, in an angry and sarcastic voice, 'don't you think your report will write itself?'

4.2 Directives and the Reasons that Back Them

How can indirect speech acts help to explain the range of directive acts performable with imperatives? The key to my response is the idea that, in directing someone to ψ , one normally also intends for them to ψ for a specific reason, and one communicates

this reason to them in some way.²¹ Although we have some methods of making these reasons for complying explicit, we often communicate them in the form of indirect speech acts. Different kinds of directives can be individuated in terms of the different reasons for complying that are offered in order to back them in this way.

To see this idea in action, consider again the intuitive distinction between requests and commands. In addressing (51) to a friend who is helping with a renovation, I intend my desire for a hammer to be reason enough for my friend to comply.

(51) Pass me the hammer.

Let us suppose that I lack any authority over my friend, and we are in a highly cooperative context with aligned interests, so they will be likely to assume that my directive is backed by nothing more sinister than my desire for a hammer. Contrast this case with one in which my boss addresses (49) to me, or in which a mugger utters (50) while aiming a gun in my direction.

(49) Get me that report.

(50) Put your hands up.

In each of these cases, I am being given clear reasons to act that go beyond the desires of the speaker. There is a clear indirect conditional threat in the case of (50): if I don't put my hands up, the mugger will (or at least might) shoot me. The reason backing (49) could be described either as a threat of this kind—if I don't give my boss the report, he might fire me or give me a bad performance review—or as an appeal to my boss's institutional authority over me.

So here is an empirical hypothesis: requests are directives backed by expressions of the speaker's desire for the addressee to comply; commands are directives backed by either (i) an appeal to the speaker's authority, or (ii) a conditional threat to frustrate their preferences in some way unless they comply. When one requests or commands with an unadorned imperative, these backing speech acts are implicit and indirect. Adding 'please' to an imperative is a conventional way of forcing a request reading, whereas one can force a command reading by adding an 'or else'-clause that makes the threatened consequences explicit.²²

²¹The idea of individuating different kinds of speech acts in terms of the reasons backing them is originally due to Schiffer (1972, ch.4), although he does not formulate the view in terms of indirect speech acts.

²²Strictly speaking, it may be better to say that sentences of the form ' ϕ or else $\triangleright\psi$ ' are used to indirectly command by directly spelling out one's conditional threat. At least, this would make sense if threatening-sounding imperative-declarative disjunctions are best understood as negative indicative conditionals in disguise, as some authors have argued they should be.

Given this way of taxonomizing, we should expect there to be as many species of directive acts as there are reasons for intending someone to comply with a directive. Take the case of non-disinterested advice, as in (56).

(56) Don't quit your job to become a mime!

Plausibly, what makes this directive an instance of advice is that it is backed by the speaker's belief that complying would be in the addressee's best interest, or would satisfy their preferences. Indeed, this sort of advice reading can be forced by making this belief explicit with a follow-up utterance. The speaker could follow up (56) with 'Your job is so cushy, and everybody hates mimes', for example.

What other kinds of reasons can back directives? Mirroring the category of commands that are backed by implied conditional threats, we can posit a category of directives backed by implied conditional promises or bribes. We could call these directives 'inducements'. Take, for example, (60), uttered by a parent to a child, while dangling a chocolate bar over their head.

(60) Clean your room!

This line of thought suggests an approach to taxonomizing speech acts that is more principled than the Austinian method of listing off English verbs that can be used to report speech and then formulating a post-hoc list of felicity conditions to go with them. Each genus of illocutionary act is defined by a distinctive type of m-intended response. Species within the genus may be further distinguished by the reasons on the basis of which speakers intend their addressees to have this response.

Why would directive acts be always, or at least normally, backed by reasons that may be communicated indirectly in this way? An answer to this question follows from the first principles of Intention-Based Semantics. In directing you to ψ , I intend for you to form an intention to ψ , partly on the basis of recognizing that I have this intention. But merely recognizing that I intend you to make a plan will almost never give you enough reason, on its own, to form this plan. Given that I'm being rational, then, I will normally have to communicate some further reason for you to comply. Given that you expect me to be rational, you will seek to discern this further reason. Given the power of the human capacity for mindreading, I see no problem with assuming that we routinely communicate reasons to one another in this way, often without making them explicit. And indeed, it is difficult to imagine a realistic and felicitous use of an imperative wherein the speaker's intended reason(s) for the addressee to comply wouldn't be relatively easy to recognize.

4.3 Making as If to Direct

What, then, of non-directive uses of imperatives, such as wishes and disinterested advice? These, on my view, are cases in which one *makes as if* to direct in order to perform only the indirect speech act that would normally be used to communicate a reason for compliance.

Take disinterested advice, such as the B's utterance, addressed to a stranger, A, in (55).

- (55) A: Excuse me, how do I get to San Francisco?
B: Take the train that leaves from over there in 10 minutes.

This looks a lot like a case of non-disinterested, directive advice, in which one directs someone to take the train, and implicates that doing so would be a good way to get to San Francisco. But in this case, B is merely *making as if* to direct A to take the train. B's real aim is to perform an indirect assertoric speech act—i.e., to answer A's question.

Wish uses can be understood in an analogous way.

- (54) Get well soon!

The speech act performed with (54) is a lot like a request, wherein S directs A to get well and implicates that this is what S desires. But here, since it is obvious that A lacks agency with respect to whether they will get well, and S knows this, it is clear that S is merely making as if to direct A to get well in order to indirectly express their desire. It the latter expression of desire that is the real point. Given that this sort of indirect speech is commonplace in the assertoric domain, it shouldn't surprise us that it turns up here, too.

Condoravdi and Lauer (2012) argue it is not directive force but the expression of a kind of desire—specifically, the speaker's "effective preference"—that is semantically encoded in imperatives. Inverting my explanation, they argue that it is directive force that is communicated only indirectly when using imperatives, and only in special circumstances in which it is clear that the speaker is expressing their own effective preference in order to influence the hearer's effective preference (2012, 48). As evidence for this claim, they argue that it is always infelicitous to produce an utterance of the form '! ϕ , but I don't want ϕ '. Their data include the following:

- (61) # Call him at home! I don't want you to but he is fine with that.
(62) A: How do I get into the building?
B: # I don't want you to but just go through this door.

I think that these data are questionable for reasons other than their form, however. Given the view I've defended, clarifying that one doesn't desire the hearer to comply will sound sensible only if one is relatively clearly communicating another reason to comply—a reason that overrides the speaker's desires. Condoravdi and Lauer's examples don't have these features, but it is not hard to come up with examples that do. The following variations on Condoravdi and Lauer's (62) strike me as perfectly fine, for example.

- (63) Go through this door. It's not what I want you to do, but it's the only way for you to get what you want.
- (64) Go through this door. It's for the best, even if it isn't what I want.
- (65) Go through this door. I don't want you to, but what I want doesn't matter in this case.

These cases work because the speaker makes their intended, non-request reading clear after the fact, and so the speaker's desire is no part of the reason for which they intend the hearer to comply. On the other hand, the following sequence sounds very bad to me.

- (66) # Open this door. It's not that I'm trying to get you to open it; I just want you to.

Of course, it is sometimes possible to produce an utterance of this form in order to force a non-directive reading. E.g.:

- (67) Get well soon! ...Not that I'm trying to get you to do anything; I just want you to get well.

But this is hardly surprising, given that, on my analysis, the speaker of such an utterance is merely making as if to direct their hearer to get well in order to indirectly express their desire. The follow-up utterance merely clarifies this fact.

Finally, what about weak readings of imperatives, such as (57)?

- (57) Have a cookie.

Can we understand these as indirect speech acts performed by making as if to perform directives? It is plausible that we can. By analogy, consider utterances of declaratives that function in ways that are weaker than assertion because of special features of the context in which they're used. Suppose Dr. House and his team are in the early stages of a differential diagnosis. At this point in the process, they're

brainstorming a list of ideas about what could be causing the patient's neurological symptoms. They'll start narrowing down the list afterward. The following dialogue unfolds.

- (68) Thirteen: A brain tumor is restricting blood flow to her hypothalamus.
Kutner: She fell and hit her head.
Taub: She has lupus.
Thirteen: She was exposed to a toxin.

Given the context, these utterances won't naturally be understood as assertions—i.e., as attempts to add their content to the addressees' beliefs. Even Stalnaker's relatively theory of assertion, on which it is an attempt to add what is asserted to the set of propositions that is presupposed or mutually accepted for the purposes of the conversation (1978; 2014), does not correctly characterize this conversation. The point of these utterances is to make their contents at least temporarily compatible with what the speakers believe or accept—or at least to bring to salience possibilities that are already compatible with what they believe or accept—so that they can then reason about which of these options best explain the patient's symptoms. We can see this from the fact that there is nothing incoherent about Thirteen suggesting two incompatible diagnoses, for example. After all: the current conversational goal is merely to get some live options on the table before deciding which one is right.

This should sound a lot like weak uses of imperatives, whose aim, I have argued, is not for the addressee to form an intention to comply, but merely for the addressee to make complying compatible with their plans, if only temporarily, for the purposes of engaging in unobstructed practical reasoning. In much the same way that brainstorming contexts force weak readings of declaratives, weak readings of imperatives pop up in cases when it's just not plausible to understand the speaker as issuing a directive. Politeness forbids forcing cookies on one's guests, for example, and so it would be unnatural, in the absence of clues that the speaker is being highly impolite, or even threatening, to understand (57) as a directive. This isn't a full explanation of how weak readings of either declaratives or imperatives arise, but I think that an explanation along these lines is at least plausible.²³

Conclusions

I have argued that imperatives semantically encode directive force, and that directive force should be understood as the attempt to get one's addressee to form an intention. I have shown how to implement this idea in a formal-semantic theory, on

²³For a similar, pragmatic explanation of weak readings of imperatives, see (von Stechow and Iatridou, ms, §4.2).

which the semantic value of an imperative clause is the kind of psychological effect that a speaker would intend to produce when issuing a literal and unembedded utterance of it. The resulting theory predicts a wide range of data, including facts about how imperatives combine with declaratives in conjunctions, disjunctions, and conditionals, as well as facts about inferences and infelicities involving imperatives that have eluded most other theories. I have also argued that this semantics is compatible with the illocutionary diversity of imperatives, provided that we adopt a theory of indirect speech acts in the vein of Grice's theory of conversational implicature.

The theory I have outlined here predicts our evidence about how imperatives work at least as well as any other theory of imperatives on the market. But the theory also goes some way toward explaining *why* imperatives work the way that they do by grounding their semantic properties in a well motivated picture of human psychology. By spelling this out, I hope to have shown that Intention-Based Semantics is of interest for more than programmatic and philosophical reasons. It can also help us to formulate independently motivated explanations of nuanced semantic phenomena.

References

- Bach, K. and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, Mass.
- Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press, Cambridge, Mass.
- Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell.
- Cantwell, J. (2008). Changing the modal context. *Theoria*, 74:331–351.
- Charlow, N. (2010). Restricting and embedding imperatives. In Aloni, M., Bastiaanse, H., de Jager, T., and Schulz, K., editors, *Logic, Language, and Meaning: Selected Papers from the 17th Amsterdam Colloquium*. Springer.
- Charlow, N. (2013). What we know and what to do. *Synthese*, 190(12):2291–2323.
- Charlow, N. (2014). Logic and semantics for imperatives. *Journal of Philosophical Logic*, 43:617–664.
- Charlow, N. (2017). Clause-type, force, and normative judgment in the semantics of imperatives. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford, UK.
- Clark-Younger, H. (2012). Is imperative inference. In MacLaurin, J., editor, *Rationis Defensor: Essays in Honour of Colin Cheyne*, pages 275–292. Springer.
- Condoravdi, C. and Lauer, S. (2012). Imperatives: Meaning and illocutionary force. In Pinõn, C., editor, *Empirical Issues in Syntax and Semantics 9: Papers From the Colloque de Syntaxe et Sémantique à Paris 2011*, pages 37–58.
- Davidson, D. (1979). Moods and performances. In Margalit, A., editor, *Meaning and Use*. D. Reidel, Dordrecht.
- Frege, G. (1894). review of ‘philosophie der arithmetik. logische und psychologische untersuchungen’ by edmund husserl. *Zeitschrift für Philosophie und philosophische Kritik*, 103:313–332.
- Fusco, M. (2015). Deontic modality and the semantics of choice. *Philosophers’ Imprint*, 15(28):1–27.
- Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3):377–388.
- Grice, H. P. (1968). Utterer’s meaning, sentence-meaning, and word-meaning. *Foundations of Language*, 4(3):225–242.
- Grice, H. P. (1969). Utterer’s meaning and intention. *The Philosophical Review*, 78(2):147–177.
- Grice, H. P. (1971). Intention and uncertainty. *Proceedings of the British Academy*, 57:263–279.
- Grice, P. (1989). *Studies in the Way of Words*. Harvard University Press, Cambridge, Mass.
- Han, C. (1998). *The Structure and Interpretation of Imperatives: Mood and Force in Universal Grammar*. PhD thesis, University of Pennsylvania.
- Harris, D. W. (2014). *Speech Act Theoretic Semantics*. PhD Dissertation, City University of New York Graduate Center.
- Harris, D. W. (MS). Semantics without semantic content. *Unpublished Manuscript*.
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press.
- Holton, R. (2011). *Willing, Wanting, Waiting*. Oxford University Press.
- Husserl, E. (1975). *Logische Untersuchungen. Erster Band: Prolegomena zur reinen Logik*, volume XVIII of *Husserliana*. Nijhoff, The Hague.
- Kaufmann, M. (2012). *Interpreting Imperatives*. Springer.
- Kaufmann, M. (2013). Embedded imperatives in colloquial german. *Language*, 89(3):619–637.
- Kaufmann, M. (2016). Fine-tuning natural language imperatives. *Journal of Logic and Computation*, Advance Access.
- Kaufmann, S. and Schwager, M. (2009). A unified analysis of conditional imperatives. In Cormany, E., Ito, S., and Lutz, D., editors, *Proceedings of SALT 19*, pages 239–256, Columbus, OH. Ohio State University.
- Kolodny, N. and Brunero, J. (2013). Instrumental rationality. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. URL = <<http://plato.stanford.edu/archives/spr2016/entries/rationality-instrumental/>>.
- Kolodny, N. and MacFarlane, J. (2010). Ifs and oughts. *The Journal of Philosophy*, 107(3):115–142.
- Murray, S. E. (2014). Varieties of update. *Semantics and Pragmatics*, 7(2):1–53.
- Murray, S. E. and Starr, W. B. (MS). The structure of communicative acts. *Unpublished Manuscript*, Online at <<https://www.dropbox.com/s/ctf8stqvp1rdqs/The-Structure-of-Communicative-Acts.pdf?dl=1>>.
- Parsons, J. (2012). Cognitivism about imperatives. *Analysis*, 72:49–54.
- Parsons, J. (2013). Command and consequence. *Philosophical Studies*, 164(1):61–92.
- Parsons, J. (2014). The transformational approach to imperative consequence. *Unpublished Manuscript*, Online at <<http://www.joshparsons.net/draft/imp-lation/imp-lation.pdf>>.
- Pelletier, F. J., Elio, R., and Hanson, P. (2008). Is logic all in our heads? from naturalism to psychologism. *Studia Logica*, 86:1–65.
- Portner, P. (2004). The semantics of imperatives within a theory of clause-types. In Watanabe, K. and Young,

- R., editors, *Proceedings of SALT 14*. CLC Publications.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15:351–383.
- Portner, P. (2012). Permission and choice. In Grewendorf, G. and Zimmermann, T., editors, *Discourse and Grammar: From Sentence Types to Lexical Categories*. Mouton de Gruyter.
- Portner, P. (2017). Commitment to priorities. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford, UK.
- Roberts, C. (2004). Context in dynamic interpretation. In Horn, L. and Ward, G., editors, *The Handbook of Pragmatics*, pages 197–220. Blackwell.
- Roberts, C. (2015). Conditional plans and imperatives: A semantics and pragmatics for imperative mood. In Brochhagen, T., Roelofsen, F., and Theiler, N., editors, *Proceedings of the 20th Amsterdam Colloquium*, pages 353–362.
- Roberts, C. (2017). Speech acts in discourse context. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*, Oxford, UK. Oxford University Press.
- Schiffer, S. (1972). *Meaning*. Oxford University Press, Oxford.
- Schiffer, S. (1982). Intention-based semantics. *Notre Dame Journal of Formal Logic*, 23(2):119–156.
- Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and Semantics 9*, pages 315–332. Academic Press, New York.
- Stalnaker, R. (2014). *Context*. Context and Content. Oxford University Press, Oxford.
- Starr, W. (ms). A preference semantics for imperatives. *Unpublished Manuscript*.
- Starr, W. B. (2010). *Conditionals, Meaning and Mood*. PhD thesis, Rutgers University.
- Starr, W. B. (2014). Mood, force, and truth. *Protosociology*, 31:160–180.
- von Fintel, K. (ms). How to do conditional things with words. *Unpublished Manuscript*.
- von Fintel, K. and Iatridou, S. (ms). A modest proposal for the meaning of imperatives. In Arregui, A., Rivero, M., and Salanova, A. P., editors, *Modality Across Semantic Categories*. Oxford University Press.
- Vranas, P. (2008). New foundations for imperative logic i: Logical connectives, consistency, and quantifiers. *Noûs*, 42:529–572.
- Vranas, P. (2010). In defense of imperative inference. *Journal of Philosophical Logic*, 39:59–71.
- Vranas, P. (2011). New foundations for imperative logic: Pure imperative inference. *Mind*, 120:369–446.
- Vranas, P. (2016). New foundations for imperative logic iii: A general definition of argument validity. *Synthese*, 193:1703–1753.
- Weisberg, M. (2007). Three kinds of idealization. *Journal of Philosophy*, 104(12):639–659.
- Yalcin, S. (2012). Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160.
- Zanuttini, R., Pak, M., and Portner, P. (2012). A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects. *Natural Language & Linguistic Theory*, 30(4):1231–1274.