

Imperatives and Intention-Based Semantics

Daniel W. Harris^{*}

July 9, 2017

Abstract

I argue that the semantic values of declarative and imperative clauses are the kinds of mental states that they are used literally to provoke in addressees: declaratives encode beliefs and imperatives encode intentions. 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. 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 Intention-Based Semantics—the research program, founded by Grice, that seeks to offer comprehensive psychological explanations of semantic and pragmatic phenomena.

1 Imperatives

My aim is to defend a theory of the semantics and pragmatics of imperative clauses that accounts for a range of puzzling phenomena about how imperatives embed, how they feature in inference, and why they can be used to perform a wide range of

^{*}I have received helpful feedback on earlier versions of some of the ideas presented here from Justin Bledin, Nate Charlow, Kai von Fintel, Daniel Fogal, Magdalena Kaufmann, Ernie Lepore, Stephen Neale, Jim Pryor, Craige Roberts, and Will Starr. 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. This draft is not final, and I would appreciate feedback, which can be emailed to me at danielwharris@gmail.com.

speech acts. My central thesis is that a clause semantically encodes the kind of mental state that it is used to produce when used literally and unembedded: declarative encode beliefs and imperatives encode intentions.¹

My main goal is to articulate my positive view and to show how it makes sense of a wide range of data. However, I will periodically compare my own theory to what I take to be the three best-worked competitor theories. First is the idea that whereas a declarative clause encodes a proposition, an imperative encodes a property—intuitively, the property that an addressee would have to realize in order to comply.² Second is the idea that imperatives encode modal propositions—that they are, in effect, equivalent to deontic modals that can only be used performatively.³ Third is a range of dynamic-semantic proposals—the idea that imperatives encode context-change potentials.⁴ None of these three theories accounts for all of the phenomena that I will try to explain here. My defense of this claim is scattered throughout what follows, but here is a brief preview of the data I will tackle and the problems they pose to alternative theories.

First, there is the fact that imperatives can be combined with declaratives in various ways—by conjunction and disjunction, as a conditional, or in both ways at once:

- (1) Make us omelettes or I'll get us some bagels.
- (2) Help me if you can.
- (3) If the manuscript is finished, send it off and I will open a bottle of wine to celebrate.

In §4, I canvass some reasons for thinking that this data makes trouble for both the view that imperatives encode modal propositions and the view that they encode properties. My solution is to think of complex, mixed-mood sentences like (1)–(3) as encoding complex mental states that can be recursively specified in terms of the beliefs and intentions encoded by their clausal parts.

Second, imperatives can partake in seemingly valid arguments:

- (4) Bake me a cake.
If you don't preheat the oven, you're not going to bake that cake!
So, preheat the oven!

¹The main precursor to my account is Charlow (2014), though the two accounts differ in a variety of ways that I will point out along the way.

²von Stechow (2004); Hausser (1980); Portner (2004, 2007, 2012, 2017); Roberts (2015, 2017); Zanuttini et al. (2012).

³Han (1998); Kaufmann (2012, 2016); Kaufmann and Schwager (2009); Lewis (1975).

⁴Condoravdi and Lauer (2012); Murray (2014); Murray and Starr (2017, MS); Starr (2016, ms, 2010).

Data about inference patterns of this kind are normally counted among the primary data of semantic theories. However, the phenomenon of imperative inference is puzzling because validity is often understood to be a matter of truth preservation, but imperatives seemingly aren't truth apt. The puzzle is particularly deep because imperatives diverge from ordinary declaratives in the inference patterns they partake in. For example, Ross-paradoxical inferences seem bad, seemingly invalidating disjunction introduction:

- (5) Give me a hug!
So, give me a hug or spit in my face!

And free-choice inferences have valid-sounding readings:

- (6) Have an IPA or a stout.
Have an IPA.

Many theorists, including those who take imperatives to encode properties, have not had anything to say about this data.⁵ In §5, I show how to predict these inference patterns with a formally precise version of the following claim: an argument will strike us as valid just in case (it is clear that) any rationally coherent mind that exemplifies the semantic values of the premises exemplifies the semantic value of the conclusion. (I remain ambivalent about whether this relation should be thought of as logical consequence.)

Third, imperatives can be used to perform a range of speech acts. Take (7):

- (7) Stay until midnight.

Addressed by a manager to a subordinate who doesn't want to stay late, (7) might be uttered as a command. Coming from a friend who has to stay and wants company, (7) would sound like a polite request. Addressed to someone who (the speaker knows) wants to stay late but isn't sure if they may, (7) would sound like an offer of permission or acquiescence. Requests and commands are both directive acts, which is to say that a speaker who performs one of these acts won't have achieved their aim unless their addressee forms an intention to comply. But acts of permitting and acquiescing aren't directives: the speaker is merely making a new option available, not necessarily trying to get their addressee to take it. What explains this range of uses that imperatives exhibit? In §6, I argue that all of these uses involve indirect speech acts. Imperatives encode neutral directive force. To perform a more specific kind of directive act, such as a command or a request, requires indirectly communicating

⁵The two main exceptions are Charlow (2014) and Starr (ms).

a particular sort of reason for the addressee to comply. To use an imperative non-directively is to make as if to direct someone to do something in order to indirectly communicate a reason why they might want to comply. I argue that this account explains the data just as well as those of property theorists and modal theorists, and that it does a better job than any of the dynamic-semantic accounts have managed.

2 Intention-Based Semantics

My theory of imperatives is intended to be a case study in *Intention-Based Semantics* (IBS), which is the research program, founded by Grice, that seeks to give comprehensive explanations of semantic and pragmatic phenomena in terms of the underlying psychological capacities involved in being a rational communicator.⁶

One component of IBS is a theory of communication. To perform a communicative act is to make an utterance with an *m-intention* (a.k.a. ‘meaning intention’ or ‘communicative intention’)—an intention to change the mind of one’s addressee, in part by revealing one’s intention to do so.⁷ Communicative acts of different kinds are distinguished by the kinds of psychological effects they are aimed at producing—a.k.a., their ‘m-intended effects’. Whereas asserting *p* is a matter of m-intending for one’s addressee to believe *p*, directing someone to φ is a matter of m-intending for one’s addressee to form an intention to φ (Grice, 1957, 1968, 1969).⁸ Communication is thus understood as an exercise of our cognitive capacity for mindreading.

A second component of IBS is the ambition to understand the semantic properties of linguistic expressions in terms of the roles they play in communication—and so, indirectly, in terms of underlying psychological capacities. For example, the role of a sentence’s meaning is to provide hearers with evidence of the intentions of a speaker who utters it (provided that the speaker is being literal). Since interpreting a communicative act is a matter of identifying the speaker’s m-intention, and m-intentions vary according to their m-intended effects, the role of a sentence’s meaning is to tell the hearer what kind of effect the speaker is trying to have on them. Semantics, on this view, is the study of the systems by means of which we offer richly structured evidence of the contents of our communicative intentions

⁶Some influential contributions to IBS include Bach and Harnish (1979); Bennett (1976); Grice (1957, 1961, 1968, 1969, 1975, 1981); Loar (1981); Neale (1992, 2004, 2005); Schiffer (1972, 1981, 1982); Sperber and Wilson (1995); Stine (1978); Strawson (1964). The term ‘intention-based semantics’ is due to Schiffer (1982).

⁷Grice would have used the term ‘utterer’s meaning’ rather than ‘communicative acts’. I adopt the latter term, which is due to Bach and Harnish (1979), because it makes for more straightforward discussion of non-assertoric cases.

⁸Note that Grice uses the term ‘indicative-type utterances’ and ‘imperative-type utterances’ for what I’m calling ‘assertions’ and ‘directives’, respectively.

(Grice, 1957; Schiffer, 2003).

I take this idea quite literally here, identifying a clause's semantic value is the m-intended effect of the communicative act that one would perform with a literal, unembedded utterance of the clause. Whereas the semantic value of a declarative is a belief—e.g., 'dogs are better than cats' encodes the belief that dogs are better than cats—the semantic value of an imperative is an intention—e.g., 'dance!' encodes the intention to dance. Formally, I model these semantic values as properties of *cognitive models*, which are idealized representations of agents' mental states.⁹

Perhaps the most interesting upshot of this account is that it allows us not just to predict, but also to *explain* data about imperative inference. On the view I will offer in §5, the intuition that an inference is valid arises from our sensitivity to the coherence relations that govern beliefs and intentions. The fact that (4) strikes us as valid, for example, is the linguistic reflex of a coherence principle usually called 'strict means-end coherence': an agent who intends to ψ , and who believes that they can't *psi* without *phi*ing, should, *ceteris paribus*, form an intention to ϕ . More generally: an argument will strike us as valid when it would be redundant to address the conclusion to a psychologically coherent agent who had already gone along with the premises. The coherence principles I appeal to here are independently motivated by decades of work in epistemology and the philosophy of action. This account therefore makes good on the central promise of IBS, which is to offer deep explanations of semantic phenomena in terms of general features of the psychology of rational communicators.

A third component of intention-based semantics is a theory of indirect communication. Interpreting a speaker's intentions is a kind of mindreading, which in turn is a kind of non-monotonic inference. Thus, although sentence meaning is powerful evidence of a speaker's intentions, it is *defeasible* evidence: in some situations, nonlinguistic evidence may make it clear that a speaker means something other than, or in addition to, what the sentence they've uttered suggests. In cases of this kind, the speaker performs an indirect communicative act.¹⁰ In §6, I argue that we normally use imperatives to perform indirect acts of various kinds.

It should be clear that I am a card-carrying intention-based semanticist. However, I think that the program has traditionally suffered from a lack of detailed and fruitful integration with semantics. Indeed, I know of no attempts to make IBS motivate positive semantic proposals. This essay is a case study in how to overcome this deficit. The result, I think, is both a plausible, independently motivated account of imperatives and a new source of support for IBS.

⁹The idea of representing clausal semantic values as properties of set-theoretic models of mental states are due to Charlow (2014).

¹⁰Grice's term was 'implicature', but he discussed only assertoric cases in detail, and once again, the term 'indirect communicative act' makes it easier to discuss non-assertoric cases.

3 Cognitive Models

The central building block of my semantics will be the notion of a cognitive model, which is an idealized representation of an agent's beliefs and desires. I will take each cognitive model M to be a pair of a Hintikka-style *belief state* B_M —the set of possible worlds compatible with what the agent being modeled believes—and an analogous *intention state* I_M —the set of worlds compatible with everything the agent intends. Importantly, for reasons that I will explain momentarily, I stipulate that for any cognitive model M , the intention state I_M must be a subset of the belief state B_M .

(8) COGNITIVE MODELS

A cognitive model M is an ordered pair $\langle B_M, I_M \rangle$ such that $I_M \subseteq B_M \subseteq W$ (where W is the universe of possible worlds).

I represent beliefs and intentions as properties of cognitive models. Taking propositions to be sets of possible worlds, this can be done as follows. A belief whose content is a proposition p can be represented as $\lambda M. B_M \subseteq p$. This is a property possessed by a cognitive model whose belief state entails p . Likewise, an intention whose content is q can be represented as $\lambda M. B_M \subseteq q$.¹¹

Cognitive models are *idealized* representations of mental states in two, importantly distinct respects. First, they distort belief and intention in various ways for the sake of simplicity. A well known bug in possible-worlds models of belief, for example, is that they collapse beliefs with necessarily equivalent contents. The present model distorts intentions in the same way. Even perfectly rational human minds would not have these qualities. I raise them here in order to set them aside. A more naturalistic model would make more accurate predictions at the cost of significant added complexity.

In the terminology of McMullin (1985) and Weisberg (2007; 2013), these are *Galilean* idealizations—distortions that are provisionally introduced into a model for the purpose of making its subject matter computationally tractable. A standard example is Galileo's model of projectile motion. By abstracting away from air resistance, achieves significant simplification while allowing for predictive power in a range of circumstances. Cognitive models are useful in a similar way: they capture some phenomena related to intention and belief while ignoring others, including phenomena that give rise to hyperintensionality—that would have to be built back into the model in order to generate a wider range of accurate predictions.

¹¹We often report intentions with non-finite 'to'-clauses, as in 'Mookie intended to do the right thing'. An agent A 's intention to ψ can likewise be represented as $\lambda M. I_M \subseteq q$, where q is the proposition that A does ψ .

The second kind of idealization at work here constitute a feature rather than a bug. Cognitive models are designed to represent agents who would be ideally structurally rational, in that their beliefs and intentions would obey a variety of coherence principles that have been posited by philosophers of action and epistemologists.¹² For example, an agent's 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. Both of these principles are guaranteed in cognitive models by the fact that belief states and intention states are modeled as sets of worlds. Other coherence principles are guaranteed by the fact that a model's intention state is stipulated to be a subset of the belief state. For example, intentions should be doxastically constrained: an agent should not intend to do things that are ruled out by their beliefs. Likewise, 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.¹³ Cognitive models are designed to guarantee these norms. This means that no actual human mind is accurately represented by a cognitive model.

This is what Weisberg calls *minimalist* idealization, the aim of which is to isolate those factors in a system that play a crucial role in explaining how the target phenomenon works (Weisberg, 2007, 642). Accordingly, the justification for minimalist idealization is not mere computational convenience, and there is no pressure to eventually lift such idealizations, since factors that are idealized away do not play a significant role in predicting or explaining the target phenomenon. A standard example, which is due to Strevens (2011), involves explanations of Boyle's law wherein "theorists often introduce the assumption that gas molecules do not collide with each other" (Weisberg, 2007, 643). This false assumption is justified because the target phenomenon is the macroscopic behavior of low-pressure gases—a phenomenon to which collisions make no difference. Similarly, one goal of cognitive models is to predict the way in which human rationality, and our sensitivity to it,

¹²These coherence principles, or 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). Note that, for my purposes, it does not matter if we treat these principles as fundamental requirements or as arising indirectly from the functional role of intention, as Kolodny (2007; 2008a; 2008b) has argued. What is important for my purposes is merely that agents tend to try to adhere to these principles, that there is something normatively defective about not adhering to them, whatever that amounts to, and that we are often sensitive to our own and others' successes and failures to adhere.

¹³ 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.

manifests itself in our intuitions about valid inference. Just as idealized models low-pressure gases would not be good for studying colliding molecules, cognitive models would not be good for studying human irrationality. But that's okay, because these are not the phenomena at issue.

4 Semantics for Declarative and Imperative Clauses

The language-fragment with which I'll be dealing 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 $\triangleright\varphi$ and an arbitrary imperative sentence will be represented as $!\varphi$.¹⁴ To keep things simple, I will assume that sentence radicals have propositions (modeled as sets of worlds) as their semantic values. For convenience, I will say that the semantic value of a sentence's radical is the sentence's (propositional, or semantic) content.

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

The syntax of the fragment is given by the following definition:

- If Φ and Ψ are both declarative sentences, then not- Φ , (Φ and Ψ), (Φ or Ψ), and (If Φ, Ψ) are also declarative sentences. Nothing else is a declarative sentence.
- If Φ and Ψ are (declarative or imperative) sentences, then (Φ and Ψ) and (Φ or Ψ) are sentences.¹⁵
- If Φ is a declarative sentence and Ψ is a sentence, then (If Φ, Ψ) is a sentence.
- Nothing else is a sentence.

¹⁴Others 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.

¹⁵Note that that '!' never scopes under negation. I see no reason to think otherwise.

4.1 Atomic Sentences

I'll give the semantics of atomic sentences by mapping each one to a type of mental state, defined, as above, as property of cognitive models. The possession of this property is the 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 $\llbracket \Phi \rrbracket(M) = 1$. By extension, I will sometimes speak of an agent satisfying a sentence, meaning that their mental state exemplifies the property encoded by the sentence.

I begin with the base case of the atomic sentences—i.e., declarative and imperative clauses themselves.

(9) DECLARATIVE CLAUSES

$$\llbracket \triangleright \varphi \rrbracket = \lambda M. B_M \subseteq \llbracket \varphi \rrbracket$$

(10) IMPERATIVE CLAUSES

$$\llbracket !\varphi \rrbracket = \lambda M. I_M \subseteq \llbracket \varphi \rrbracket$$

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.

4.2 Conjunction and Disjunction

Imperatives can be both conjoined and disjoined with other imperatives.

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

(12) Take the bus or walk.

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

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

(14) 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 (13) could seemingly have gotten the same point across with (15), and one might as well utter (16) as (14).

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

(16) 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. 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.¹⁶

I have been convinced by Starr (2016; ms), however, that not all mixed conjunctions and disjunctions are disguised conditionals.¹⁷ To my ears, for example, the most prominent reading of (17) 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.

(17) 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.

(18) 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.

¹⁶I won't try to account for the conditional readings of mixed conjunctions and disjunctions here. I agree with von Fintel and Iatridou (ms) that any approach to this phenomenon will have to understand them as part of the broader category of conditional conjunction, on which see Culicover and Jackendoff (1997); Keshet (2012).

¹⁷Notice that some other constructions are likewise ambiguous between conditional-conjunction and regular conjunctive readings, e.g.: 'He touches my friend and he's going to jail.'

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 given 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:

- (19) 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? When the conjunction mixes clauses of both types, 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 recognizing their existence.

Disjunction is a bigger challenge.¹⁸ 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 may be in a state of indecision with respect to the disjuncts. The kind of mental state satisfying a disjunction would have 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 may postpone the choice between the two until later. You need only enter a state of mind that commits you to eventually choosing at least one of these two alternatives.

How must a cognitive model be structured in order to represent minds with these properties? I will consider two options, both of which will be easier to lay out with the help of the notion of a *submodel* of a cognitive model.

(20) Φ -SUPPORTING SUBMODEL

For any sentence Φ and any cognitive model M , a Φ -supporting submodel of M is a cognitive model M_Φ such that:

- (i) $\llbracket \Phi \rrbracket (M_\Phi) = 1$;
- (ii) $B_{M_\Phi} \subseteq B_M$; and
- (iii) $I_{M_\Phi} \subseteq I_M$.

A cognitive model M that satisfies a disjunction needn’t satisfy either disjunct. Rather, there must be two submodels of M , each of which satisfies a respective disjunct, and the union of these subsets must be M . Here are two candidate clauses for disjunction that both embody this idea in slightly different ways.¹⁹

¹⁸Charlow’s (2014) semantics is closest to mine of any in the literature, but as Starr (ms) points out, his theory can’t make sense of imperative-imperative or imperative-declarative disjunctions.

¹⁹I use the symbol ‘ \cup ’ for pointwise union. In general, if $A = \langle a_1, \dots, a_n \rangle$ and $B = \langle b_1, \dots, b_n \rangle$,

(21) WEAK DISJUNCTION

$$\llbracket \Phi \text{ or } \Psi \rrbracket = \lambda M . (\exists M_\Phi)(\exists M_\Psi)M_\Phi \cup M_\Psi = M$$

(22) STRONG DISJUNCTION

$$\llbracket \Phi \text{ or } \Psi \rrbracket = \lambda M . (\exists M_\Phi : I_{M_\Phi} \neq \emptyset)(\exists M_\Psi : I_{M_\Psi} \neq \emptyset)M_\Phi \cup M_\Psi = M$$

It follows from each of these two clauses that one can satisfy a disjunction by being in a state of indecision between two states of mind, each of which would satisfy one of the disjuncts. For example, suppose that I have not yet decided whether to buy IPA or stout, but I am committed to at least one of the two options. In that case, I satisfy (23) by the lights of either weak disjunction or strong disjunction.

(23) Buy IPA or buy stout.

The two clauses differ with respect to whether someone can satisfy a disjunction by satisfying only one of the two disjuncts. Suppose, for example, that I have made up my mind to buy IPA and not stout. According to the weak disjunction clause, I do satisfy (23) in this state of mind, whereas I don't according to the strong disjunction clause. According to the strong disjunction clause, satisfying a disjunction requires being in a state of indecision between the two disjuncts.

It may initially seem obvious that natural-language disjunction is weak in the sense just outlined. After all: if I perform a literal and direct speech act with (23), and, as a result, my addressee forms an intention to buy IPA, shouldn't I count as having succeeded?

However, there are some reasons to prefer strong disjunction. The motivating idea here is that, in uttering a disjunction, the speaker intends for their addressee to treat both disjuncts as live options, even if only temporarily and for the purpose of reasoning about which to settle on. In cases where an addressee immediately chooses one disjunct or the other, they won't have done as the speaker intended unless they reached this state by at least momentarily treating both disjuncts as live options and deliberating about which to settle on.

To see the plausibility of this idea, consider some of the different roles that uttering a disjunction of imperatives, such as (24), could play in a discourse.

(24) Mow the lawn or wash the car.

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 (24) in order to get them to choose at least one of these two courses of action

$A \cup B = \langle a_1 \cup b_1, \dots, a_n \cup b_n \rangle$. If M^1 and M^2 are both cognitive models, then $M^1 \cup M^2 = \langle B_{M^1} \cup B_{M^2}, I_{M^1} \cup I_{M^2} \rangle$.

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 about which to adopt. On this way of thinking about things, if you comply with my utterance of (24) 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 intentions, and then choosing between them.²⁰

Now consider a different scenario, in which B utters (24) as part of the following dialogue.

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

Here, A's utterance makes it clear A intends to wash the car. So, if disjunction were weak, then B's reply would be redundant: A would already satisfy B's whole disjunction by satisfying its first disjunct. 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 disjunction of declaratives.

- (26) 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.

There is some independent motivation for thinking that disjunction is semantically strong in the sense that I have suggesting. For example, Zimmermann (2000) and Geurts (2005) both draw on facts about free-choice phenomena and the semantics of lists to argue for accounts of disjunction that make $\varphi \vee \psi$ semantically

²⁰Note that much of this discussion assumes that the two disjuncts of a disjunction are treated as exclusive alternatives by the participants in a conversation, as they normally are. However, neither of the proposed clauses has a problem with inclusive uses of disjunction, which is just a limit case. Any cognitive model $M = M_\Phi = M_\Psi$ would satisfy the disjunction ' Φ or Ψ ' by satisfying both disjuncts.

equivalent to a corresponding conjunction of possibility modals, $\diamond\phi \wedge \diamond\psi$. If we combine this view with a relatively straightforward semantics for \diamond , then the Zimmerman/Geurts view is equivalent to STRONG DISJUNCTION (at least when it comes to declaratives).²¹ Aloni (2017) likewise uses free-choice phenomena to defend an account of disjunction that is equivalent to STRONG DISJUNCTION, so long as we restrict our attention to declaratives. My suggestion that we adopt strong disjunction can be considered a generalization of these views that handles imperatives as well as declaratives.

Another possibility is that disjunction is semantically weak, but that it often behaves in the strong way for pragmatic reasons. Simons (2005, §5.3.1) defends a view of this kind (again, designed to handle only declaratives). She argues that disjunctions are subject to a pragmatic “genuineness constraint” that both disjuncts be epistemically possible. Her main reason for thinking that this constraint is pragmatic is that it is sometimes cancellable. She borrows and the following example from Grice (1989, 45) :

I can say to my children at some stage in a treasure hunt, *The prize is either in the garden or in the attic. I know that because I know where I put it, but I'm not going to tell you.* Or I could just say (in the same situation) *The prize is either in the garden or in the attic*, and the situation would be sufficient to appraise the children of the fact that my reason for accepting the disjunction is that I know a particular disjunct to be true [and therefore the other is not an epistemic possibility].

As Simons’ interprets her genuineness constraint, each disjunct of a disjunction must be epistemically possible *for the speaker*, and so Grice’s scenario would count as canceling it. As Simons (2005, 300n42) points out, Zimmermann (2000, 271n24) also takes the epistemic possibility of disjuncts to pertain to the speaker’s information state, and so he takes examples like Grice’s to make trouble for his semantic view as well—albeit not enough trouble that he gives up the view. However, we should take care to notice that these considerations do not apply to STRONG DISJUNCTION, given my preferred interpretation as a part of intention-based semantics. On my view, a speaker who uses a disjunction literally needn’t themselves take both disjuncts

²¹For example, the following clause for \diamond would do the job:

$$(27) \llbracket \diamond\phi \rrbracket = \lambda M . B_M \cap \llbracket \phi \rrbracket \neq \emptyset$$

This resembles Veltman’s (1996) semantics for epistemic modals, but rendered in the present framework. The intuitive idea behind this clause is that the aim of uttering $\diamond\phi$ is to make the addressee’s beliefs compatible with the content of ϕ . If I tell you that it might be raining, for example, my aim is to put you into a belief state that does not rule out the possibility that it is raining. I won’t defend this view here, but I think that it is plausible enough.

to be live options. Rather, they must only intend for the *hearer* to treat both disjuncts as live options. This requirement is satisfied in Grice's scenario.

Why does it matter which account of disjunction we adopt? Although I think that the issue is of some independent interest, its significance will also re-emerge in §5. Foreshadowing that discussion: strong disjunction will be needed in order to give a semantic account of the badness of Ross's paradox and the goodness of imperative free-choice inferences. And so, if we take natural-language disjunction to be weak and so offer a pragmatic account of examples like (25) and (26), we will also be forced to give a pragmatic account of these inferential phenomenon, as, e.g., Simons (2005) does. Although I won't claim to have settled this issue here, the above considerations make me at least somewhat optimistic about STRONG DISJUNCTION.

Finally, what about mixed-clause disjunctions? Either clause will work for our purposes. In uttering (18)a with its true disjunctive sense, for example, I intend to put my addressee into a mental state of at least momentary indecision between intending to put back Waverly and believing that I will put back Naked Lunch.

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

Why try to put someone in this state of mind? A plausible reason would be 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 may, in turn, require action on their part in order to coordinate their beliefs and plans with mine. On the other hand, upon entering into the mental state that is the semantic value of (18), 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 (18)a.

4.3 Conditionals

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

(27) Attack if the weather is good.

(28) 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 (28), for example, one shouldn't attempt to buy the speaker a drink immediately, but only once one forms the belief that the bartender has come back.

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

- (29) If the bartender is back, buy me a drink and I'll find us a table.
- (30) 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; von Stechow, ms; Roberts, 2015). 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, *ceteris 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*:

- (31) MAXIMAL SUBMODEL
 A maximal Φ -satisfying submodel of a cognitive model M is a cognitive model $M_{\uparrow\Phi}$ that meets the following conditions:
- (i) $M_{\uparrow\Phi}$ is a Φ -supporting submodel of M
 - (ii) There is no Φ -supporting submodel M^* of M such that either $B_{M_{\uparrow\Phi}} \subset B_{M^*}$ or $I_{M_{\uparrow\Phi}} \subset I_{M^*}$.

With this notion in hand, we can say that a model satisfies a conditional if and only if each of its maximal antecedent-satisfying submodels satisfies its consequent.²² Formally:

²²This 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.

- (32) CONDITIONAL
 $\llbracket \text{If } \Phi, \Psi \rrbracket = \lambda M. (\forall M_{\uparrow\Phi})(\llbracket \Psi \rrbracket(M_{\uparrow\Phi}) = 1)$

The intuitive idea here is that any structurally rational state of mind that satisfied both a conditional and its antecedent would also, *ceteris paribus*, satisfy its consequent.

5 Inference and Infelicity

Some inferences involving imperatives strike us as compelling in something like the way that valid inferences do.²³

- (33) a. Take out the trash and mow the lawn!
 b. So, take out the trash! (Charlow, 2014, §2.2)
- (34) a. Attack at dawn if the weather is fine.
 b. The weather is fine.
 c. So, attack at dawn! (Parsons, 2013, 61)
- (35) (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 in order to handle cases involving imperatives. One way to appreciate the puzzle posed by these arguments is to consider the following inconsistent triad (Clark-Younger, 2012; Parsons, 2014):

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

Each of (36)i–iii has some support. The support for (36)i comes from the intuition of validity that we get from some arguments such as the above examples. 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; Starr, 2014).

²³If you doubt that imperative inference is a genuine phenomenon, see Vranas (2010) to be convinced otherwise.

Meanwhile, the support for (36)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 preservation.

Two particularly difficult data for a theory of imperative inference to explain are Ross's paradox, which raises the problem of explaining why (37)c doesn't seem to follow from (37)p, and free choice inferences, which raise the problem of explaining why (38)ci and (38)cii *do* seem to follow from (38)p (though only when (38)ci and (38)cii are being used weakly, in order to offer acquiescence or permission, and not when they are understood as directives).

- (37) ROSS'S PARADOX
(p) Post the letter.
(c) \neq Post the letter or burn the letter.
- (38) 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 predicted by classical logic, we have to get creative.

Finally, in addition to apparent instances of entailment that involve mixtures of imperatives and declaratives, there are also apparent cases of mixed imperative-declarative inconsistency (Starr, ms):

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

More generally, if it is clear that the addressee can't ψ , for whatever reason, and $[\varphi]$ is the proposition that the addressee ψ s, then it will sound inconsistent or infelicitous to utter the imperative $!\varphi$. 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.

5.1 A Theory of Imperative Inference

With the help of the semantics laid out in §4, I can explain this data by defining a relation \vdash as follows.²⁴

²⁴This definition of \vdash is closely related to a variety of dynamic and/or informational consequence relations, as found in, e.g., Bledin (2014); van Benthem (2011); Veltman (1996); Yalcin (2007). For a similar proposal that applies to inferences involving imperatives, see Starr (ms, 2010).

(41) INFERENCE

$\Phi_1, \dots, \Phi_n \vdash \Psi$ iff:

$(\forall M)$ if $\llbracket \Phi_1 \rrbracket(M) = 1, \dots, \llbracket \Phi_n \rrbracket(M) = 1$, then $\llbracket \Psi \rrbracket(M) = 1$

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) on a structurally rational mind that already satisfied 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 (33) and (34) 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 (36). 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.²⁵ Instead, it makes logic the study of how the norms governing the structural rationality of human thought get reflected in natural language. This view resolves the triad by rejecting (36)iii.

One reason not to think of \vdash as the relation of logical consequence is the desire to preserve the traditional idea that genuine logical consequence is about truth preservation.²⁶ A related reason is that this proposal would commit us to psychologism about logical consequence, according to which logic is the study of features of 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.²⁷

²⁵It is natural, on my view, to think of a sentence as a truth-bearer just in case its semantic value is a truth-bearing mental state. Declaratives are truth-apt because beliefs are truth-apt; imperatives are non-truth-apt because intentions are non-truth-apt. Formally, a sentence Φ is true iff, for any intention state I , $\llbracket \Phi \rrbracket(\langle \{ @ \}, I) = 1$, where $@$ is the actual world. A sentence Φ is truth-apt just in case its semantic value is intention-state-insensitive—i.e.: for any cognitive model M such that $\llbracket \Phi \rrbracket(M) = 1$, and for any intention state I^* , $\llbracket \Phi \rrbracket(\langle B_M, I^* \rangle) = 1$.

²⁶Of course, not everyone thinks that this traditional idea is a good one. See, e.g., Bledin (2014).

²⁷For some classic attacks on psychologism about logic, see Frege (1894); Husserl (1975). For a

So, I prefer to understand \vdash as a model of ways in which the coherence norms governing our 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 we are sensitive to the norms that are constitutive of practical rationality as they apply to both ourselves and others. 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 (36)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. 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, (35), which is an instance of \vdash .

- (35) (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 principle of strict means-end coherence. Any agent who intends to buy me an IPA, and who believes that buying me an IPA requires going to the bar, is rationally required, by dint of this principle, 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 (35); it also explains why we find such inferences compelling by appealing to an independently motivated norm of human psychology.

more recent survey on the relevant issues, see Pelletier et al. (2008).

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 and p , if $B_M \cap p = \emptyset$, then $I_M \not\sqsubseteq p$.) This property of cognitive models allows us to predict the infelicity of (39) and (40).

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

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

The two halves of (39) 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 (39) is explained. The same goes for (40), modulo the 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.

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.

(37) ROSS'S PARADOX

(p) Post the letter.

(c) $\not\equiv$ Post the letter or burn the letter.

The system presented here correctly blocks this inference, but only if we take disjunction to be strong, in the sense outlined in §4.2. If disjunction is weak, then satisfying (37)p (i.e., intending to post the letter) is a way of satisfying (37)c, and so any cognitive model that satisfies the premise of Ross's paradox also satisfies its conclusion. However, if disjunction is strong, then satisfying (37)c requires treating both posting and burning as live options—a state of mind that is not compatible with intending to do one of the two. As in our other examples, moreover, the fact that Ross's paradox is not intuitively valid can be explained by a fact about the nature of intending: if one intends to post the letter, it does not follow that both posting and

burning the letter are compatible with one's plans. Indeed, the two mental states would normally be incompatible.

Even if disjunction is semantically weak, it is clear from some of the data presented in §4.2 that we often use 'or' to convey a strong-disjunctive meaning. If that's the right way to think about disjunction, then a pragmatic account of Ross's paradox may be the right way to go as well. Roughly, the idea behind such an account would be that we reject the inference in (37) because we read the pragmatically strengthened meaning into (37)c, and this meaning doesn't strike us as following from (37)p.

Next, consider FREE CHOICE.

- (38) 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 an 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 noted earlier, free choice inferences 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 (38). Other instances of the same pattern don't look as good, seemingly because their conclusions are most naturally understood as strong, directive acts:

- (42) (p) Pay the toll or turn around!
(c) $\not\models$ 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:

- (43) Have a cookie. [offer]
(44) (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 merely attempting to make a new

practical option 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 'iφ' (Charlow, 2014; Starr, ms).²⁸ The intuitive contrast between strong and weak uses is this: whereas !φ is used to change the addressee's plans in such a way that they intend $\llbracket\varphi\rrbracket$, iφ is used to make what the addressee intends compatible with $\llbracket\varphi\rrbracket$ (cf. Charlow 2014).

- (45) ATOMIC WEAK-IMPERATIVE CLAUSES
 $\llbracket i\varphi \rrbracket = \lambda M. I_M \cap \llbracket \varphi \rrbracket \neq \emptyset$

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

- (46) $(! \varphi \text{ or } ! \psi) \vdash i \varphi$

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

Capturing free choice inferences semantically therefore depends on two controversial assumptions about the semantics–pragmatics interface—that disjunction is semantically strong, and that imperative clauses are semantically ambiguous between strong and weak readings. I have already discussed some of the complexities surrounding the former issue. I will briefly consider how weak readings of imperatives could arise pragmatically in §6. If either strong readings of disjunction or weak uses of imperatives are best explained pragmatically, then we might be able to give a pragmatic explanation of free choice inferences that is compatible with the rest of the present account. Roughly: in many circumstances we tend to either pragmatically strengthen the premise or pragmatically weaken the conclusion of free-choice inferences (or both), which leads us to interpret them as if they behave according to the semantics just presented. Although more would have to be said to justify such a view, I don't think it is unreasonable, and I won't attempt to definitively rule it out here.

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

²⁸One reason to think that weak uses are at least sometimes semantically encoded is that they are strongly preferred when an imperative is used with rising intonation (Portner, 2017).

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 $!\varphi \vdash !(\varphi \vee \psi)$ and $!(\varphi \vee \psi) \vdash !\varphi$, 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 and up.

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 to 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.

In particular, I think that this explanation of imperative inference gives my theory an advantage over dynamic-semantic theories, such as the theory of Starr (2016, ms, 2010). Starr’s is the only other available theory that predicts all of the same data about embedding and inference that I predict here,²⁹ but his theory does so by postulating that conversations revolve around baroquely structured discourse contexts with dimensions that represent interlocutors shared information, unresolved questions, and preferences. The properties of these discourse contexts are stipulated so as to predict a range of linguistic data, but Starr does not attempt to show how they can be independently motivated in terms of broader facts about the nature of communication or human psychology. So although the theory I have presented here is roughly equivalent to Starr’s with respect to the linguistic data it predicts, I take my account to be superior for reasons related to broader principles of theory choice. This advantage is a direct result of my theory’s integration with intention-based semantics.

²⁹It is worth noting, however, that Starr does not attempt to explain the data about illocutionary variability that I will consider in §6.

6 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 several species within the directive genus as well as 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 (48) to a subordinate, or someone uttering (49) while aiming a gun at their addressee.

(48) Get me that report.

(49) Put your hands up.

Requests form another subcategory of directives. Addressed to a friend who is helping with a renovation, (50) would normally be understood as a request rather than a command, for example.

(50) 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.

(51) Pass me the hammer please.

(52) 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 (53) to a friend in the hospital, for example, I can't rationally intend to change their plans, since I know (and know that they know) that they lack control over the outcome.

(53) Get well soon! (Condoravdi and Lauer, 2012, 39)

Instead, it looks like this use of (53) is merely a way of expressing a wish or preference for my friend to improve. Similarly, imperatives are sometimes used to offer what might be called ‘instructions’ or ‘disinterested advice’, as in B’s half of (54).

(54) 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 (55) to me in an urgent tone, it seems clear that their goal is to sway my action.

(55) Don’t quit your job to become a mime!

However, it is consistent with an exchange like (54) that B doesn’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 §5, including permission or invitation uses, as in (56), as well as what von Stechow and Iatridou (ms, 5) call ‘acquiescence’ uses, as in (57), and ‘indifference’ uses, as in (58).

(56) Have a cookie.

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

(58) 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 §5, 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 Stechow 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 §§6.2–6.3. First, I will say more about what I mean by indirect speech acts.

6.1 Indirect Speech Acts

Indirect speech acts, on the view that I prefer, subsume Grice's 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 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. 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 believe *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 conversational 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 and 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 the hearer can 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, 'do you think your report will write itself?'

6.2 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 account 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.³⁰ Although we have some methods for 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 (50) 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.

(50) 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 that my friend 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 (48) to me, or in which a mugger utters (49) while aiming a gun in my direction.

(48) Get me that report.

(49) 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 (49): if I don't put my hands up, the mugger will (or at least might) shoot me. The reason backing (48) 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 the addressee's 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

³⁰The idea of individuating different kinds of speech acts in terms of the reasons backing them is originally due to Strawson (1964) and Schiffer (1972, ch.4), although they do not formulate the view in terms of indirect speech acts.

signaling a request reading, whereas one can force a command reading by adding an ‘or else’-clause that makes the threatened consequences explicit.³¹

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 (55).

(55) 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 the addressee’s 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 (55) 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, (59), uttered by a parent to a child, while dangling a chocolate bar over their head.

(59) Clean your room!

This line of thought suggests a principled approach to taxonomizing speech acts. Each genus of illocutionary act is defined by a distinctive type of m-intended effect. Species within the genus may be further distinguished by the reasons on the basis of which speakers intend their addressees to have this response.³²

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

³¹Strictly 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.

³²Again, roughly this idea is due to Schiffer (1972, ch.4).

and felicitous use of an imperative wherein the speaker's intended reason(s) for the addressee to comply wouldn't be relatively easy for the addressee to recognize.

It is instructive to compare this view to the most influential alternative accounts of why imperatives can be used to perform a wide range of speech acts. According to one account, imperatives are modals that can only be used performatively, and the particular kind of speech act one performs in uttering an imperative is determined by the contextually specified ordering source relative to which this modal is interpreted (Kaufmann, 2012). When an imperative is understood as a command, its ordering source is something along the lines of *what the speaker orders*, for example, whereas when it is interpreted as a request or a wish, its ordering source is something more like *what the speaker desires*. This ordering source determines the speech act being performed by changing the ranking according to which certain possible worlds are deemed ideal, according to the speech act: to comply with a command is to bring the actual world closer to those that are ideal with respect to what the speaker orders, whereas to comply with a request is to bring the actual world closer to those that are ideal with respect to what the speaker desires. On Portner's (2007) view, by contrast, all uses of imperatives add a new action to the addressee's To-Do List, but directives of different kinds update different subsections of the To-Do List, each of which plays a role in influencing different kinds of ordering sources for subsequent deontic modals. Successfully commanding someone to ψ will tend to produce the result that they ought to ψ *in view of what the speaker authoritatively requires*, whereas successfully requesting that someone ψ will tend to produce the effect that they ought to ψ *in view of what the speaker desires*.

The roles of ordering sources in the former account and of the subsections of the To-Do List in the latter account are analogous to the role of directive-backing indirect speech act in my account. Each represents a kind of information that a hearer must infer, beyond recognizing what is encoded in a sentence, in order to correctly identify the kind of directive being performed. Moreover, in each case, the information to be inferred concerns not *what* the addressee is being directed to do but *why* they should comply. The difference lies in the nature of what the hearer is inferring. On the modal account, the speaker must infer which conversational background is the imperative's intended ordering source. On Portner's account, the hearer must infer which subsection of their To-Do list the speaker is proposing to update. On my account, the hearer must interpret the indirect speech act that is backing the speaker's directive. Moreover, in any of these cases, we should expect successful communication to affect the future direction of the conversation, including, for example, how future deontic modals are interpreted. In general, however information about my desires becomes clear, for example, the interpretation of subsequent deontic modals whose ordering source is *what I desire* will be altered as a result. So, if I indirectly communicate that I desire for you to buy me a drink by

requesting for you to do so, this will normally make subsequent utterances of (60) come out true.

(60) (In view of what I desire,) you ought to buy me a drink.

Although this isn't quite a form of modal subordination, as, e.g., Kaufmann (2012) and Roberts (2017) would have it, it is what we should expect to happen for principled pragmatic reasons.

6.3 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 someone to do something in order to indirectly perform the 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 (54).

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

This is similar to 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 not actually trying to get A to form an intention to take the train, and this is obvious to A. The real purpose of B's utterance is merely to impart the information that taking the train would be a good way to get to San Francisco. In other words: 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 that answers A's question.

Wish uses can be understood in an analogous way. Suppose Steve addresses (53) to Ann:

(53) Get well soon!

The speech act performed with (53) is a lot like a request, wherein Steve directs Ann to get well and indirectly informs her that this is what he desires. But here, since it is obvious that Ann lacks agency with respect to whether she will get well, and Steve knows this, it is clear that Steve is merely making as if to direct Ann to get well in order to indirectly express his desire. It is the latter expression of desire that is Steve's 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 don’t think that this point generalizes. Given the view I’ve defended, we should expect that following an imperative with a clarification to the effect that one doesn’t want the addressee to comply will often sound bad, except in situations in which an alternative, desire-overriding reason for the addressee to comply is communicated in some way. Condoravdi and Lauer’s examples don’t have this feature, 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 felicitous, 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, clarifying that their 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. In short: I don't think that Moore-paradox-style data should carry much weight here.

Finally, what about weak uses of imperatives, such as (56)?

(56) 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, I think, although I won't try to give a detailed account. 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 theory of assertion, on which it is a proposal to make a proposition 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 (56) 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.³³

Conclusions

I have argued that imperatives semantically encode directive force, and that a directive 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 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 spelled out 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 the psychology of rational communicators. By spelling this out, I hope to have shown that Intention-Based Semantics is of interest for more than programmatic and foundational reasons. It can also help us to formulate independently motivated explanations of nuanced semantic phenomena.

³³For a similar, pragmatic explanation of weak readings of imperatives, see von Stechow and Iatridou (ms, §4.2).

References

- Aloni, M. (2017). Epistemic and deontic modals in state-based semantics. *presentation slides from Inquisitiveness Below and Beyond the Sentence Boundary*.
- Bach, K. and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, Mass.
- Bennett, J. (1976). *Linguistic Behaviour*. Cambridge University Press, Cambridge, UK.
- Bledin, J. (2014). Logic informed. *Mind*, 123(490):277–316.
- 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. (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.
- Culicover, P. W. and Jackendoff, R. (1997). Semantic subordination despite syntactic coordination. *Linguistic Inquiry*, 28(2):195–217.
- Davidson, D. (1979). Moods and performances. In Margalit, A., editor, *Meaning and Use*. D. Reidel, Dordrecht.
- 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.
- Frege, G. (1894). review of ‘philosophie der arithmetik. logische und psychologische untersuchungen’ by edmund husserl. *Zeitschrift für Philosophie und philosophische Kritik*, 103:313–332.
- Geurts, B. (2005). Entertaining alternatives: Disjunctions as modals. *Natural Language Semantics*, 13:383–410.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3):377–388.
- Grice, H. P. (1961). The causal theory of perception. *Proceedings of the Aristotelian Society, Supplementary Volume*, 35:121–153.
- 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. (1975). Logic and conversation. In Cole, P., editor, *Syntax and Semantics 3: Pragmatics*. Academic Press.
- Grice, P. (1981). Presupposition and conversational implicature. In Cole, P., editor, *Radical Pragmatics*. Academic Press, New York.
- 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.
- Hausser, R. (1980). Surface compositionality and the semantics of mood. In Searle, J., Kiefer, F., and Bierwisch, M., editors, *Speech Act Theory and Pragmatics*, pages 71–96. Reidel, Dordrecht.
- 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. (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.
- Keshet, E. (2012). Focus on conditional conjunction. *Journal of Semantics*, 30(2):211–256.
- Kolodny, N. (2007). How does coherence matter? *Proceedings of the Aristotelian Society*, 107(3):229–263.
- Kolodny, N. (2008a). The myth of practical consistency. *European Journal of Philosophy*, 16(3):366–402.
- Kolodny, N. (2008b). Why be disposed to be coherent? *Ethics*, 118:437–463.
- 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.
- Lewis, D. K. (1975). A problem about permission. In et al., E. S., editor, *Essays in Honour of Jakko Hintikka*. D. Reidel, Dordrecht, The Netherlands.

- Loar, B. (1981). *Mind and Meaning*. Cambridge University Press, Cambridge, UK.
- McMullin, E. (1985). Galilean idealization. *Studies in History and Philosophy of Science*, 16(3):247–273.
- Murray, S. E. (2014). Varieties of update. *Semantics and Pragmatics*, 7(2):1–53.
- Murray, S. E. and Starr, W. B. (2017). Force and conversational states. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford, UK.
- Murray, S. E. and Starr, W. B. (MS). The structure of communicative acts. *Unpublished Manuscript*, Online at <<https://www.dropbox.com/s/ctf8stqvpl1rdqs/The-Structure-of-Communicative-Acts.pdf?dl=1>>.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15:509–559.
- Neale, S. (2004). This, that, and the other. In Bezuidenhout, A. and Reimer, M., editors, *Descriptions and Beyond*, pages 68–182. Oxford University Press.
- Neale, S. (2005). Pragmatism and binding. In Szabó, Z. G., editor, *Semantics versus Pragmatics*, pages 165–285. Oxford University Press.
- 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-xlation/imp-xlation.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. (2015). Conditional plans and imperatives: A semantics and pragmatics for imperative mood. In Brochhagen, T., Roelofsen, E., 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. (1981). Indexicals and the theory of reference. *Synthese*, 49(1):43–100.
- Schiffer, S. (1982). Intention-based semantics. *Notre Dame Journal of Formal Logic*, 23(2):119–156.
- Schiffer, S. (2003). *The Things We Mean*. Oxford University Press.
- Simons, M. (2005). Dividing things up: The semantics of ‘or’ and the modal/‘or’ interaction. *Natural Language Semantics*, 13:271–316.
- Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell, Oxford, 2 edition.
- 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. (2016). Conjoined imperatives and declaratives. In *Proceedings from Sinn and Bedeutung 21*. University of Edinburgh.
- 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.
- Stine, G. (1978). Meaning other than what we say and referring. *Philosophical Studies*, 33(4):319–337.
- Strawson, P. F. (1964). Intention and convention in speech acts. *The Philosophical Review*, 73(4):439–460.
- Strevens, M. (2011). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.
- van Benthem, J. (2011). *Logical Dynamics of Information and Interaction*. Cambridge University Press, Cambridge, UK.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261.
- Vranas, P. (2010). In defense of imperative inference. *Journal of Philosophical Logic*, 39:59–71.
- Weisberg, M. (2007). Three kinds of idealization. *Journal of Philosophy*, 104(12):639–659.
- Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford University Press, Oxford and New York.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.
- 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.
- Zimmermann, T. E. (2000). Free choice disjunction and epistemic possibility. *Natural Language Semantics*, 8:255–290.