



Imperative inference and practical rationality

Daniel W. Harris¹ 

Accepted: 19 June 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract Some arguments include imperative clauses. For example: ‘Buy me a drink; you can’t buy me that drink unless you go to the bar; so, go to the bar!’ How should we build a logic that predicts which of these arguments are good? Because imperatives aren’t truth apt and so don’t stand in relations of truth preservation, this technical question gives rise to a foundational one: What would be the subject matter of this logic? I argue that declaratives are used to produce beliefs, imperatives are used to produce intentions, and beliefs and intentions are subject to rational requirements. An argument will strike us as valid when anyone whose mental state satisfies the premises is rationally required to satisfy the conclusion. For example, the above argument reflects the principle that it is irrational not to intend what one takes to be the necessary means to one’s intended ends. I argue that all intuitively good patterns of imperative inference can be explained using off-the-shelf formulations of our rational requirements. I then develop a formal-semantic theory embodying this view that predicts a range of data, including free-choice effects and Ross’s paradox. The resulting theory shows one way that our aspirations to rational agency can be discerned in the patterns of our speech, and is a case study in how the philosophy of language and the philosophy of action can be mutually illuminating.

I have received helpful feedback on earlier versions of this paper from Chris Barker, Bob Beddor, Justin Bledin, Lucas Champollion, Nate Charlow, Josh Dever, Kai von Fintel, Daniel Fogal, Magdalena Kaufmann, Ernie Lepore, Stephen Neale, Paul Portner, Jim Pryor, Craige Roberts, Will Starr, Elmar Unnsteinsson, Alex Worsnip, Seth Yalcin, several anonymous referees, and audiences at the University of Pennsylvania, Hunter College, the New York Philosophy of Language Workshop, NYU, the Institute of Philosophy in London, PhLiP 4 in Tarrytown, and the Chapel Hill Normativity Workshop.

✉ Daniel W. Harris
danielwharris@gmail.com

¹ Department of Philosophy, Hunter College, The City University of New York, 695 Park Ave., New York, NY 10065, USA

Keywords Imperatives · Imperative inference · Imperative logic · Rational requirements · Practical rationality · Speech acts · Semantics · Metasemantics

1 Two challenges posed by imperative inference

We meet up at a bar. The first thing I say to you is (1):

(1) Buy me a drink.

You nod, but don't do anything else. After a few seconds, I say (2):

(2) You can't buy me that drink unless you go to the bar.

'Right!', you say, but don't move. Growing frustrated, I utter (3):

(3) So go to the bar!

In this scenario, I have made an argument—one whose second premise and conclusion were initially left implicit, but which I slowly spelled out. Here is the argument in full:

(4) Buy me a drink.

You won't buy me that drink unless you go to the bar.

So, go to the bar!

This is, in some sense, a good argument.¹ By contrast, (5) is a non-sequitur:

(5) Buy me a drink.

You won't buy me that drink unless you go to the bar.

So, go to the beach!

Both (4) and (5) are what I will call 'imperative arguments'—arguments that include imperative clauses in their premises, their conclusion, or both. The fact that some but not all imperative arguments strike us as good raises two challenges, one technical and one foundational.

The technical challenge is to define a consequence relation that predicts which imperative arguments are good. This is a challenge because imperative arguments apparently don't obey classical logic. For example, Ross (1941) pointed out that disjunction-introduction fails when imperatives are involved:

(6) ROSS'S PARADOX

Post the letter.

≠ So, post the letter or burn the letter.

On the other hand, imperative free choice arguments, such as (7), seem valid if their conclusions are interpreted as permissions, acquiescences, invitations, or other

¹ For a defense of the idea that there are good imperative arguments, see Vranas (2010).

“weak” speech acts (but not if their conclusions are interpreted as commands or other “strong” or “directive” acts).²

- (7) IMPERATIVE FREE CHOICE
 Have some tea or have some coffee.
 Have some tea.

Imperative arguments may include only imperative clauses, as in (7), or they may involve a mix of imperative and declarative clauses, as in (4). Likewise, imperatives can strike us as inconsistent with either other imperatives or declaratives:

- (8) Buy me a drink!
 # Don't buy me a drink!
- (9) Unicorns don't exist.
 # Bring me a unicorn!
 Starr (2020)

We therefore can't quarantine imperatives from declaratives by specifying one logic for the former and another for the latter. To predict these data, we need a consequence relation that can handle clauses of both kinds. The technical challenge is to formulate this relation.

The foundational challenge is to say what the subject matter of this consequence relation would be. The standard view is that a consequence relation is a theory of which inference patterns are necessarily truth preserving. But this account doesn't work for imperative arguments because imperatives seemingly aren't truth apt.³ So if we manage to design a consequence relation that separates good from bad imperative inferences, what will we be giving a theory of?

My strategy will be to tackle the foundational challenge first and use the resulting account to motivate a formal semantics and consequence relation that meet the technical challenge. My answer to the foundational challenge builds on a picture of the semantics and pragmatics of imperative and declarative clauses that is inspired, in its broad strokes, by Grice (1968, 1969). On this view, humans communicate by overtly attempting to change their addressees' minds in various ways, and communication succeeds if the addressee recognizes how the speaker intended to change their mind. An assertoric communicative act is an attempt to get the addressee to form a belief, and a directive communicative act is an attempt get them to form an intention. Language facilitates communication by allowing us to encode richly structured (but defeasible) evidence of the sorts of mental states that we

² On imperative free choice, see Aloni (2007). On the distinction between strong and weak uses of imperatives, see von Fintel and Iatridou (2017); Starr (2020); Charlow (2014).

³ Some have held that imperatives *are* truth apt, contrary to appearances (e.g. Lewis (1970) and Kaufmann (2012)). I lack the space to say why I think they are wrong, but see Charlow (2014, 2018) and Starr (2014).

intend to produce by speaking. In particular, declarative clauses encode beliefs, and are paradigmatically used to perform assertoric acts, and imperative clauses encode intentions, and are paradigmatically used to perform directive acts.⁴ For example, if I am speaking literally, then my point in uttering (1) is to get you to form an intention to buy me a drink. By semantically representing the intention to buy me a drink, the sentence gives you evidence that producing this mental state in you is my aim in speaking.⁵

Thinking of the semantic value of a clause as the mental state that it is used to produce opens the possibility of explaining properties of the clause in terms of properties of the mental state. This will be my strategy. Specifically, I will argue that our validity and consistency judgments reflect our sensitivity to the rational requirements that govern beliefs and intentions. Rational requirements are a major topic in the philosophy of action. Here I will argue that they do important explanatory work in the philosophy of language as well.

Consider the example with which I began. I explain the fact that (4) is an intuitively good argument by appeal the fact that we are subject to a rational requirement of means-end coherence, which can be stated as follows:⁶

(10) MEANS-END COHERENCE

An agent who intends to ϕ and believes that they must ψ in order to ϕ , but who fails to intend to ψ , is in this respect irrational.

If I have had my intended effect by uttering the two premises of (4), then you intend to buy me a drink and believe that you must go to the bar in order to do so. If you are in these states but don't intend to go to the bar, then you fail to be means-end coherent. This is what explains our intuition that (4) is a good argument: it would be irrational to be in the states that are semantically encoded by the premises without also being in the state that is semantically encoded by the conclusion.

I will say that a certain mental state m satisfies a sentence s if m is the mental state that a speaker intends to produce in an addressee when using the sentence literally. And I will say that an agent satisfies s if they are in a mental state that

⁴ Some have argued that directive force is not semantically encoded in imperative clauses. One commonly given reason is that we sometimes perform non-directive speech acts using imperatives. For example, 'have a cookie' might be used to invite or permit someone to have a cookie; 'get well soon' might be used to wish someone well; and 'take the A train' might be used to give instructions for going uptown. None of these speech acts would naturally be described as directives, nor need they count as failures if the addressee fails to form an intention to comply. I think that these data are compatible with my view, but an adequate explanation of why would not fit into this paper, and so will have to await another occasion. However, I will discuss "weak" uses of imperatives, such as invitations and permissions below (see in particular footnote 15). I will also briefly discuss wish uses in footnote 10. I think that instruction uses are best understood as indirect speech acts in which the speaker makes as if to issue a directive in order to indirectly communicate information about how one would comply with it, but I won't try to defend this view here.

⁵ Charlow (2014) defends the view that declaratives denote beliefs and imperatives denote plans (which I take to be interchangeable with intentions). His formal semantics is a major influence on the one that I will defend in Sect. 6. It is unclear how much of my foundational account Charlow would accept.

⁶ For defenses of variations on this principle, see Hill (1973); Bratman (1987, 31–2); Broome (1999; 2013, §9.4) Schroeder (2004); and Holton (2008, 51).

satisfies s . Put in these terms, my claim is that an argument from $p_1 \dots p_n$ to c will strike us as valid if any agent who satisfies $p_1 \dots p_n$ is rationally required to satisfy c as well.

Imperative inconsistencies can be explained in a similar way. For example, the felt inconsistency of (8) and (9) can be explained in terms of our sensitivity to the following two rational requirements, respectively:

- (11) CONSISTENCY OF INTENTIONS⁷
 An agent who intends to ϕ and who intends to ψ , but whose beliefs rule out the possibility that they will both ϕ and ψ , is in this respect irrational.
- (12) DOXASTIC REQUIREMENT ON INTENDING⁸
 An agent who intends to ϕ but whose beliefs rule out the possibility that they will ϕ is in this respect irrational.

In order to satisfy both imperatives in (8), you would have to both intend to buy the speaker a drink and intend not to buy them a drink—a pair of clearly inconsistent intentions. In order to satisfy both sentences in (9), you would have to both believe that there are no unicorns to be had and intend to bring the speaker a unicorn—a violation of the doxastic requirement on intending. Putting aside practical limitations on our ability to appreciate the consequences of rational requirements, if it would violate a rational requirement to be in a mental state that jointly satisfies a set of sentences, they will strike us as inconsistent. On the view that I will defend, all of our intuitions about validity and inconsistency can be explained in terms of rational requirements in the way I have just illustrated.

My plan is as follows. After fleshing out my response to the foundational challenge in Sects. 2–4, I will use it to motivate an answer to the technical challenge in Sects. 5–7. Along the way I will point out some of the ways that competing accounts of the semantics and logic of imperatives fare worse than my account, although a detailed comparison of my account with others is beyond what I can accomplish here.

2 Imperatives and intentions

Some rival accounts of imperatives have been formulated in terms of mental states other than intentions. One tradition takes imperatives to be devices for expressing desires about the addressee's actions (Searle 1969; Stokke 2014). Others have linked the meanings of imperatives to the preferences that they are used to express (Condoravdi and Lauer 2012) or the preferences that they are aimed at getting others to adopt (Starr 2010, 2020). This raises the question of why I have formulated

⁷ For variations on this principle, see Bratman (1987, 31); Broome (2013, 136, 156–7); Holton (2011, 41–2).

⁸ Some have claimed that belief that one will ϕ is constitutive of intending to ϕ , so that it is impossible, and not merely irrational, to violate the doxastic constraint (Grice 1971; Harman 1976; Audi 1973; Anscombe 1963; Neale 2016). I follow a range of others in thinking that it is possible but irrational to violate the constraint (Bratman 1987; Broome 1999; Holton 2011).

my account in terms of intentions rather than one of these other states. The answer is that my aim is to explain the logical properties of imperatives and declaratives in terms of their connections to the rationality of certain states of mind, and only intentions have the right rational profile to do this explaining.⁹

Desires can't do the work of intentions, for example, in part because they are not subject to rational requirements. If I believe that a doughnut is an unhealthy breakfast, then it is irrational for me to both intend to eat a doughnut and intend to have a healthy breakfast. Still, I might *want* a doughnut and also *want* to eat something healthy. Although I would be better off if I didn't have jointly unsatisfiable desires, it is inevitable that creatures like us sometimes do, and this alone does not render us irrational. It is only when we turn our desires into intentions that the threat of genuine irrationality looms. Rationality therefore normally requires that we plan in such a way as to overcome some of our conflicting desires, but not necessarily that we eliminate them. Given that my account explains our intuitions about validity and consistency in terms of our sensitivity to rational requirements, it is therefore important that I take imperatives to be semantically linked with intentions rather than desires.

For related reasons, preferences can't do the work of intentions in the kind of explanation that I wish to give. It can be perfectly rational to have preferences about matters over which one has no control. I would prefer that the sun not go supernova next week, for example. And although I believe that I have no control over whether I will have nightmares about being unprepared to teach my classes, I have a strong preference not to have them. By contrast, it is a corollary of the irrationality of intending what we believe we can't do that it is never rational to intend something if one has no control over whether it will happen. This is why it would normally be infelicitous to tell someone to stop having nightmares after acknowledging that they lack control over whether they will. Theories according to which imperatives are devices for expressing or coordinating on preferences lack an explanation of these facts, and thereby put too much distance between the meanings of imperatives and the actions that we use them to elicit.¹⁰

⁹ My claims about rational requirements draws heavily on work by Bratman (1987, 1999, 2007), Broome (1999, 2013), Ferrero (2009), Holton (2008), (2011), and others. It is worth noting that even if rational requirements are mere side effects of our responsiveness to reasons, as Kolodny (2008a, 2008b) argues, or even if they are merely rough generalizations about how the pressure to be rational plays out, as Fogal (2020) argues, this does not prevent them (or the facts giving rise to them) from playing the role that I take them to play here.

¹⁰ Condoravdi and Lauer (2012) actually embrace this consequence, and argue that "wish" uses of imperatives (as when we tell a sick person to 'get well soon') are semantically basic, expressing the speaker's preferences, with directive uses resulting only in contexts where the speaker's preferences are understood to entail obligations for the addressee. However, most imperatives can't be felicitously used to express mere wishes (e.g. 'get tenure!'; 'win the lottery!'), and so it seems likely that these uses are idiosyncratic and do not reflect semantic properties of imperatives in general (cf. Kaufmann 2019, 341n39).

3 Logic and rationality

My explanation of our logical intuitions extends to arguments consisting entirely of declarative clauses, so long as we adopt some version of the requirement that a rational agent ought to believe what they take to be the consequences of their beliefs. For example, Broome (2013, 157) defends a version of the following principle:

- (13) MODUS PONENS REQUIREMENT
If an agent believes p , believes that if p then q , cares whether q , but fails to believe q , then the agent in this respect fails to be rational.

This is a relatively weak requirement. It does not require that an agent believe all of the logical consequences of their beliefs, or even all of what they take to be the consequences of their beliefs. (Broome adds in the clause about caring whether q in response to Harman's (1986, 12) point that it is not irrational to avoid cluttering one's mind with the trivial consequences of one's beliefs.) Still, if we assume (i) that we are at least somewhat sensitive to facts about what follows from our beliefs, and (ii) that anyone who bothers to consult their intuitions about an argument cares about its conclusion in the relevant sense, then Broome's weak requirement is sufficient to explain our intuitions about the validity of declarative arguments.

Equipped with (13), my account aims to explain our intuitions not just about imperative arguments but also about the usual declarative-only arguments as well. Am I therefore seeking to give a theory of the nature of logical consequence that competes with more traditional views? No. Such a position would involve rejecting the idea that validity is a matter of truth preservation and embracing psychologism about logic—two positions that would fly in the face of a venerable tradition. More importantly, this position would not account for the validity of arguments that our cognitive limitations prevent us from recognizing. I prefer to think of what I have sketched as part of a psychological explanation of our intuitions about validity and consistency. Our sensitivity to coherence requirements is part of a causal explanation for our judgments about validity and consistency in natural language. We will likely want to explain the nature of rational requirements themselves partly in terms of an objective notion of truth preservation defined over the contents of thoughts. But truth preservation alone cannot yield all of our rational requirements, since it cannot explain requirements on mental states that are not truth apt. Our rational requirements, and so our intuitions about validity, therefore outstrip what logic alone requires.

4 Complex sentences and complex states of mind

The account that I have sketched so far needs to be complicated in order to make sense of the ways that imperatives can combine with one another and with declaratives in complex sentences. As (6) and (7) illustrate, my account needs to make sense of disjunctions of imperatives. The same goes for conjunctions of

imperatives, as in (14), and conditional imperatives, as in (15). Imperatives can also be conjoined and disjoined to declaratives, and these imperative–declarative mixes can themselves be conditionalized, as in (16).

- (14) Take out the trash and mow the lawn.
So, take out the trash. (Charlow 2014, §2.2)
- (15) Stop if there is a stop sign.
There is a stop sign.
So, stop! (Vranas 2010, 59)
- (16) If we only have an hour, walk the dog and I'll feed the baby.
We only have an hour.
So walk the dog!

My account says that we intuit these to be good arguments because anyone with mental states that satisfy their premises is rationally required to be in a mental state that satisfies their conclusions. But this raises the question of what kind of mental states satisfy conjunctions and disjunctions of imperatives, conditional imperatives, and complex combinations of imperatives and declaratives. Since the imperative arguments that violate classical logic all involve either disjunctions or conditionals, there is reason to think that our way of handling these constructions will be particularly important.

I will argue that each of these constructions is satisfied by a kind of complex mental state that arises in practical reasoning and that has been discussed by philosophers of action. In particular, disjunctions of imperatives are satisfied by disjunctive intentions, which have been most carefully theorized by Holton (2008, 2011), and conditional imperatives are satisfied by conditional intentions, whose most detailed treatment is due to Ferrero (2009). Imperative–declarative mixes are satisfied by combinations of intentions and beliefs that arise most typically when planning collective actions, as outlined by Bratman (2014). My plan in this section is to show that these off-the-shelf accounts make excellent sense of the data.

I begin with the simplest cases involving conjunction. As (14) and (17) suggest, the usual rules of conjunction introduction and elimination apply when imperatives are involved.

- (17) Give me five dollars.
Take me to the airport.
So, give me five dollars and take me to the airport.

There is nothing mysterious about these cases. To satisfy each of the premises of (17), one must have an intention. To satisfy its conclusion, one must have both intentions. In general, satisfying a conjunction requires being simultaneously in the states of mind that satisfy its conjuncts, whether these states of mind are beliefs, intentions, or a mix of the two.

I take conditional imperatives to be satisfied by conditional intentions. For example, the aim of a literal utterance of ‘stop if there is a stop sign’ is a conditional

intention to stop if there is a stop sign. Whereas an unconditional intention is a commitment to act at the appropriate time, a conditional intention commits the agent to act at the appropriate time only given that they take the condition to have been met (Ferrero 2009).

Conditional intentions arise in the context of contingency planning. Our beliefs and intentions are not monolithic representations of how things are and what to do. We also consider ways in which our unsettled epistemic issues might come to be resolved and how we will settle our practical issues as a result. Suppose that I intend to watch the World Cup Final at a venue that is sympathetic to whoever is playing against France. At the moment, all I know is that this could be England or Croatia. So I form two conditional intentions: to go to the Croatian Cultural Center if Croatia wins the semifinal, and to go to the Churchill Pub if England wins. Reasoning about contingencies resembles practical reasoning more generally, and our contingency plans are subject to rational requirements that are analogous to those that govern our unconditional beliefs and plans (Holton 2008; Ferrero 2009). For example, my conditional intention to go to the Churchill Pub if England is playing, together with my belief that one normally needs a reservation to watch big games there, puts rational pressure on me to form an intention to make a reservation if England is playing. This is a conditionalized manifestation of means-end coherence—one that explains why the following argument is compelling:

- (18) Go to the Churchill Pub for the final if England is playing.
 If they're playing, you'll need to have made a reservation to go there.
 So, if England is playing, make a reservation.

Likewise, it is irrational to intend both to go to the Churchill Pub if England is playing and to avoid the Churchill Pub if England is playing, or to intend to go to the pub if they are playing and to believe that I won't be able to go if they are playing. These facts about the coherence requirements on conditional intentions and beliefs are what explain why the following two pairs seem inconsistent:

- (19) If England is playing, go to the Churchill Pub.
 # Don't go to the Churchill Pub if England is playing.
- (20) You won't go to the Churchill Pub if England is playing.
 # If England is playing, go to the Churchill Pub.

In some cases, we form an intention to do something knowing that we might fail, and engage in further contingency planning to prepare for the possibility of failure. Bratman (1987, 39–40) imagines a scenario in which he awakes to find a fallen tree blocking his driveway. He forms an intention to move the tree himself, but believes he will probably fail, and so makes a backup plan to call an arborist if he can't move it himself. Bratman's intention to move the tree himself, together with his conditional intention to get an arborist to do it if he can't, is the kind of mental state that one would intend to produce with literal utterances of the following two sentences:

- (21) Move the tree yourself.
If you can't move it yourself, hire an arborist.

So-called contrary-to-duty imperatives, such as (21), have been one source of complication in logics of deontic modals and imperatives since they were pointed out by Chisholm (1963).

Planning for contingencies requires keeping track of how our practical options depend on the ways that our information might change. As my last example demonstrates, we sometimes face uncertainty by ranking our practical options, with the top-ranked option being what we intend and lower-ranked options being our backup plans. On other occasions, we formulate a range of alternative plans to go with different ways that things might turn out, but without ranking these options with respect to each other. Holton (2008, 465) illustrates this by modifying Bratman's example about the tree.¹¹

Three possibilities have occurred to you. You might lever [the tree] out of the way with a crowbar (though you are not sure how much force you can exert with that). You might saw it into manageable pieces with a chainsaw (though you are not sure that you will be able to get the chainsaw started). Or you might put a rope round it and drag it out of the way with your car (though you are not sure, given the way that the car is penned in, that you will be able to manoeuvre it into a position from which this would work). ... In the end...you walk down to the shed and load up the wheelbarrow with your biggest crowbar and the chainsaw and a rope—all in preparation for a morning attempt to move it, one way or another, yourself.

In this case, Holton (2008, §6) argues, you don't all-out intend to take any one of the three options that you're considering. In particular, you aren't yet decided about whether to try the crowbar, the chainsaw, or the rope first, and so there is no good way to describe what you intend solely by attributing an unconditional Plan A together with a conditional Plan B. In this case, what you have is a *disjunctive intention*—an intention to move the tree with the crowbar, to cut it with the chainsaw, or to pull it by tying it to your car (Holton 2008, 43–4). This, I claim, is the state of mind that one intends to put someone in by uttering a disjunction of imperatives, such as (22).

- (22) Lever the tree with a crowbar, (or) cut it up with a chainsaw, or drag it with a rope tied to your car.

If an utterance of (22) successfully puts the addressee into a state of disjunctively intending as above, they may quickly move beyond that state. They may decide to eliminate one of the options, or to rank one of the options above the others, forming an all-out intention to do it and demoting the other options to backup plans. This should not be surprising given the role of intentions (including disjunctive intentions) as inputs to practical reasoning. My claim is that satisfying a disjunction

¹¹ I have left out some details from Holton's example that aren't relevant for my purposes.

of imperatives requires having a corresponding disjunctive intention at least provisionally and for the purpose of further practical reasoning.

In addition to disjunctive intentions, Holton posits a category of *partial intentions* to describe situations like his tree case. In general, when one has a disjunctive intention to ϕ or ψ , partial intention is the attitude that one bears both to ϕ and to ψ . According to Holton, each of one's partial intentions is paired with a corresponding 'partial belief' about what one can accomplish.¹² In Holton's tree case, for example, you have one partial intention to use the crowbar, together with a partial belief that you are strong enough; you have another partial intention to use the chainsaw, together with a partial belief that you will manage to start it; and you have a third partial intention to drag the tree with your rope, together with a partial belief that you can do that. An agent has a partial intention to ϕ , according to Holton (2008, 41), when they have narrowed what they intend down to a collection of alternatives, and ϕ ing is one of these alternatives. Analogously, Holton thinks of a partial belief as one of a collection of live epistemic alternatives: to partially believe p is to treat p as a live possibility which has at least one non- p possibility as a live alternative (Holton 2008, 39).

I will also take disjunctions of declaratives to be aimed at producing disjunctive beliefs, which I will think of as collections of alternative partial beliefs. For example, you satisfy 'either the Mets will win the world series or the Yankees will win', if you rule out the possibility that anyone other than the Mets or Yankees will win, while leaving each of those possibilities open—even if only just for long enough to reason about which option to rule out.

One reason to take this sort of line on disjunction is that a more traditional approach won't work in a semantic framework that pairs sentences with the mental states that satisfy them. For example, it won't do to say that a disjunction is satisfied just in case at least one of its disjuncts is satisfied. One can satisfy a disjunction by being in a state of indecision with respect to the disjuncts. You can satisfy 'get pilsner or get IPA' without having chosen which of the two beers to get, for example, and you can satisfy 'England will win or Croatia will win' without having settled on which of these two possibilities is actual.¹³

More controversially, it follows from my approach that satisfying one disjunct but not the other is never enough to satisfy the whole disjunction. If you intend to get pilsner and aren't treating IPA as a live option, then your mental state does not satisfy 'get pilsner or get IPA'. Likewise, if you believe that the Mets will win the world series, and don't consider the Yankees winning to be a live possibility, then you do not satisfy 'the Mets will win the World Series or the Yankees will win'.

¹² A confusing-terminology warning: (i) Bratman uses 'partial plan' to describe a plan whose implementation details aren't yet worked out, but this is not what Holton means by 'partial intention'. (ii) It is common to use 'partial belief' to discuss credence (a.k.a. subjective probability), but this isn't what Holton means by 'partial belief' (Holton 2008, §3).

¹³ One weakness of Charlow's (2014) precursor to my view is that he gives the semantics for a disjunction, Φ or Ψ , by defining it as equivalent to $\neg(\neg\Phi \text{ and } \neg\Psi)$. There are two problems with this. First, it incorrectly entails that satisfying a disjunction requires satisfying at least one of its disjuncts. Second, as Starr (2020, footnote 16) has pointed out, imperatives apparently don't scope under negation, and so it is unclear what a relevant negated conjunction of negated imperatives is supposed to mean.

Disjunctions, I will provisionally claim, are always and not merely sometimes devices of indecision, which require their disjuncts to be treated as live alternatives. This is a controversial view, though there is some precedent for it in the semantics literature on disjunction.¹⁴

An advantage of this way of understanding disjunction is that it allows for a simple explanation of Ross's paradox and imperative free choice. First consider Ross's paradox.

- (6) ROSS'S PARADOX
 Post the letter.
 ≠ So, post the letter or burn the letter.

My view once again explains the badness of (6) in terms of a congruent rational requirement on intentions. Someone who intends to post the letter is not rationally required to disjunctively intend to post the letter or burn it. On the contrary, the two states are rationally incompatible, given the belief that one can't both post and burn the letter. Having a disjunctive intention requires treating both disjuncts as live options, which is not compatible with having settled on one of them.

The second puzzling datum is imperative free choice—the intuitive goodness of some cases of inferring one disjunct from an imperative disjunction:

- (7) IMPERATIVE FREE CHOICE
 Have some tea or have some coffee.
 Have some tea.

It is important that this pattern looks good only if its conclusion is given a weak (permissive) reading rather than a strong (directive) reading. Weak readings are particularly clear when an imperative is uttered in response to a question about what the addressee is permitted to do, or about whether the speaker would object to some action, for example:

- (23) A: It's getting warm. Can I open the window?
 B: Sure! Go ahead. Open it! (von Fintel and Iatridou 2017, 291)

von Fintel and Iatridou (2017) point out that weak uses include the act of expressing indifference as to the addressee's actions, and that weak uses, unlike strong uses, aren't inconsistent with their negations.

- (24) Go left! Go right! I don't care. (von Fintel and Iatridou 2017, 291)

Notably, it looks like we don't change the meaning of (24) if we convert the first two sentences to a disjunction:

- (25) Go left or go right! I don't care.

¹⁴ In particular, see Zimmermann (2000); Geurts (2005); Simons (2005); Aloni (2003, 2007). I will discuss this part of my view further, and show how to weaken the assumption if desired, in footnote 29.

Similarly, notice that whereas (26) can have either a weak reading (e.g., if addressed to a guest) or a strong reading (e.g., if addressed by an interrogator to a suspect), the second half of (27) has the effect of disambiguating the first half to a weak reading.

- (26) Take a seat.
 (27) Take a seat.
 ...or not, it's up to you.

These data suggest that weak imperatives are satisfied by partial intentions, in Holton's sense, rather than by all-out intentions. The aim of a weak use of an imperative is to make a course of action available to the addressee, but to leave the choice of whether to take it up to them. There is nothing irrational about having conflicting partial intentions, and so we should expect (24) to seem consistent and to be roughly synonymous with (25). And although a disjunctive intention is incompatible with all-out intending either disjunct, it rationally requires partially intending each disjunct. This explains free-choice effects, and why they arise only with weak conclusions.¹⁵

Finally, consider conjunctions and disjunctions that mix imperatives and declaratives. For example:

- (28) Buy me a drink and I'll find us a table.
 (29) Put back Waverly or I'll put back Naked Lunch. (Starr 2020)

There are significant consequences if we take mixed coordinations like these at face value. We will need semantic clauses for 'and' and 'or' that can compose with declaratives, imperatives, or both at once. Since there is widespread consensus that only expressions with matching semantic types can be coordinated, mixed coordinations give us the basis of an argument that imperatives and declaratives have the same semantic type. This is a problem for semantic theories like Portner's (2004), which distinguishes the meanings and speech-act potentials of declaratives and imperatives by saying that the former denote propositions and the latter denote properties.¹⁶ Since mixed coordinations can also be conditionalized, we will also need a single semantic clause for the conditional that is neutral about whether its consequent is an imperative, a declarative, or a mixture of both. This is a problem for theories that give the semantics of conditional imperatives and indicative conditionals by means of separate clauses.¹⁷

¹⁵ There remains the question of whether the distinction between weak and strong uses of imperatives should be understood as a semantic ambiguity or as the result of some pragmatic effect. I will offer a semantic clause for weak imperatives in Sect. 6, but one could also take strong uses as basic and think of weak uses as arising from a kind of pragmatic weakening (von Stechow and Iatridou 2017). I won't pursue this idea here, except to say that if it is right that weak uses are derived pragmatically, then my theory predicts that imperative free-choice effects are also a pragmatic phenomenon.

¹⁶ See also Hausser (1980); Portner (2007, 2012); von Stechow and Iatridou (2017); Roberts (2018). The same problem arises for Barker (2012), who takes imperatives to denote actions, which he models as relations on worlds.

¹⁷ For example, Charlow (2014).

Some have held that when imperatives appear to be coordinated with declaratives, they have conditional rather than normal imperative meanings.¹⁸ For example, the point of uttering ‘buy me a drink and I’ll pay you back’ could be gotten across equally well by uttering ‘if you buy me a drink, I’ll pay you back’. And the point of uttering ‘pay up or I’ll squeal’ could be gotten across just as well by uttering, ‘if you don’t pay up, I’ll squeal’.

However, I have been convinced by Starr (2020) that not all mixed conjunctions or disjunctions have conditional meanings. The most prominent reading of (28) is one on which my finding a table is not meant to be understood as conditional on your buying me a drink. One could get the same point across by saying ‘I’ll find us a table. Buy me a drink’, for example. Similarly, Starr argues that (29) needn’t have a conditional meaning, but can naturally be used to propose alternative courses of action for the addressee or speaker. It could be paraphrased by saying, ‘Put back Naked Lunch; alternatively, I’ll put back Waverly.’ Although I agree that an account of mixed coordinations with conditional meanings is needed,¹⁹ I will be concerned only with genuinely conjunctive and disjunctive readings here.

The key to understanding mixed coordinations is to note a pattern of severe restrictions on their felicitous use that (to my knowledge) has not been previously discussed in the literature. A mixed disjunction can be felicitous only if its declarative disjunct predicates a future action of either the addressee or a third-party, for example.

- (30) # Buy me a drink or the square root of 2 is greater than 1.
 (31) i. Put back Waverly or I [will/can/should/am going to] put back Naked Lunch.
 ii. # Put back Waverly or I [am putting back/was putting back/have put back] Naked Lunch.
 (32) iii. Put back Waverly or Kelsey [will/can/should/is going to] put back Naked Lunch.
 iv. # Put back Waverly or Kelsey [was putting/is putting/has put] back Naked Lunch.

Similarly, mixed conjunctions sound bad when their declarative conjunct fails to predicate a future action of an agent.

- (33) ?? Buy me a drink and I have found us a table.
 (34) ?? Buy me a drink and I have a younger brother.

The pattern can be summed up by saying that mixed coordinations are felicitous typically only when they propose one or more ways of carrying out a shared plan. Mixed conjunctions are normally used to propose a pair of actions for the addressee and the speaker (or a third party) that constitute a plan for how to pursue a shared goal. For example, (28) is a proposal about how to divide up the task of getting a

¹⁸ See, e.g. Culicover and Jackendoff (1997); Charlow (2014); von Stechow and Iatridou (2017).

¹⁹ For some attempts, see Klindinst and Rothschild (2012); Keshet (2012); Keshet and Medeiros (2019); Kaufmann and Schwager (2009).

drink together. Mixed disjunctions propose alternative plans for accomplishing a shared goal, one that would be pursued by the addressee and another that would be pursued by the speaker (or another agent). For example, (29) proposes two alternative ways of staying under budget at the book sale. When these conditions aren't met by a mixed coordination, we tend to read it either as defective or as a conditional in disguise.

There is a relatively clear pragmatic explanation for these data. In order for an utterance of a conjunction to be felicitous, it must be possible to understand both conjuncts as partial answers to a single, contextually salient question. An utterance of a disjunction will be felicitous only if the disjuncts can be understood as alternative answers to a contextually salient question.²⁰ When used literally, imperatives give answers to questions about what to do (or to plan on doing). Stereotypically, they are answers to the question of what the addressee should do. But they may also be partial answers to the question of what a group of agents that includes the addressee should do. The latter are the occasions on which a mixed coordinations can be felicitous, because these are the occasions on which multiple agents' actions can be partial answers to the question of what to do.

I have been speaking somewhat loosely. On my view, clauses don't denote actions but mental states. Why should your *intention* to buy me a beer be a partial answer to the same question as your *belief* that I will find us a table? And why should one person's *intention* to put back Waverly be an alternative answer to the same question as their *belief* that someone else will put back Naked Lunch?

The answer lies in the roles that these intentions and beliefs play in joint planning. Earlier I focused on the role that Bratman and later philosophers of action assign to intentions in the *intrapersonal* coordination of one's own plans and actions. But Bratman has also emphasized the role of intentions in *interpersonal* coordination. Our intentions are stable-enough commitments that they allow others to plan around our future actions. In order to play this role, agents engaged in joint planning must have plans that fit together. They must have "meshing subplans", in Bratman's (1992, 2014) terminology. If our joint plan is to buy five books, for example, we can't each intend to buy three. But this sort of coordination won't work unless each of us has accurate beliefs about the other's intentions. As Bratman (1987, §3.4.2) puts it, A's intention and B's expectations about A's intentions are "mutually supportive" when they are conducive to successful coordination. This is why an intention to put back Waverly and a belief that one's counterpart will put back Naked Lunch can be relevant alternatives: they are alternative contributions that the addressee of (29) can make to an interpersonally coordinated subplan of their and the speaker's shared plan to buy five books.

So far, I have tried to show that we can explain much of the peculiar semantic and inferential behavior of imperatives by taking their semantic function to be the

²⁰ On the idea that assertions in general must be answers to a contextually salient question in order to be felicitous, see Roberts (2012). My claim about conjunctions follows straightforwardly. On the idea that felicitous disjunctions must present alternative answers to a contextually salient question, see Simons (2000, §2.3.2). I will not attempt to incorporate contextual questions into the formal model that I construct in the next several sections, and so these constraints on felicity will remain informal here.

production of intentions, and by taking the semantic function of declaratives to be the production of beliefs. I turn now to the problem of how to embody the foregoing ideas in a formal-semantic theory.

5 Cognitive models

My semantics is built around cognitive models, which are idealized representations of agents' beliefs and intentions. Each cognitive model M will determine a 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.²¹ An important aspect of any cognitive model M will be that $I_M \subseteq B_M$. This is to say that cognitive models represent agents as intending only outcomes that are compatible with their beliefs.

Cognitive models are idealized in that they are designed to represent ideally rational agents, who obey all of the rational requirements that I canvassed earlier.²² Any agent whose mental state could be accurately represented by a cognitive model would be guaranteed to have consistent beliefs and intentions, for example, because I model belief states and intention states as non-empty sets of worlds. Other coherence principles are guaranteed by the fact that a model's intention state is always a subset of its belief state. This guarantees that any agent who could be represented by a cognitive model would be means-end coherent and have doxastically constrained intentions.²³ Cognitive models are designed to guarantee that these requirements are met. This presumably means that no actual human mind is accurately represented by a cognitive model. This is a feature of the model rather than a bug: Since my goal is to show how our sensitivity to rational requirements shows up in how we reason with language, I am idealizing away from cases involving irrationality. In the real world, we sometimes find ourselves in irrational states of mind and we sometimes reason poorly, but cognitive models ignore these failures.²⁴

In Sects. 2–4, I showed why a model of the human mind that is adequate for present purposes needs to capture the ways that agents plan for contingencies by distinguishing and ranking different epistemic alternatives and the practical

²¹ This way of modeling beliefs is originally due to Hintikka (1962). Possible-worlds models have been used to model intentions by computer scientists working in the Belief-Desire-Intention (BDI) tradition—e.g., Cohen and Levesque (1990). The most direct precursor to cognitive models are Charlow's (2014) "representors", although they also bear some resemblance to Starr's (2020) "conversational states".

²² They are also idealized in a second sense, which is that they oversimplify some aspects of belief and intention, such as by representing agents as believing and intending all necessary truths—a well-known issue with possible-worlds models.

²³ We can model the state of believing that p is required for q as $\lambda M.(B_M \cap p) \subseteq q$. This is too strong, in that it says nothing about the p happening before q . In order to capture the weaker interpretation, we would need to lift another idealization in cognitive models by adding a temporal dimension.

²⁴ Here I am inspired by Weisberg's (2007, 642) discussion of "minimalist idealizations," which work by misrepresenting a system in a way that isolates just those features that play a crucial role in explaining the target phenomenon.

alternatives that depend on them. Cognitive models incorporate this structure as follows:

(35) COGNITIVE MODELS

A cognitive model M is a pair $\langle \Omega_M, \leq_M \rangle$ such that:
 $\Omega_M \subseteq \{ \langle \beta, \iota \rangle : \emptyset \subset \iota \subseteq \beta \subseteq W \}$ ($\Omega_M \neq \emptyset$)
 \leq_M is a preorder on Ω_M

Suppose that M represents the mental state of a certain agent—call her ‘Mae’. Ω_M is the set of Mae’s *options*—roughly, her contingency plans. Each option represents one way that, Mae anticipates, her beliefs might change (β) along with a representation of how her plans would change as a result (ι). Mae might have formulated a range of options to herself, and she might have ranked some of them over others—a ranking represented by the preorder \leq_M . We can think of the set of Mae’s top-ranked options as her Plan A, which I annotate as Ω_M^A :

(36) PLAN A

$\Omega_M^A = \{ \omega \in \Omega_M : (\forall \omega' \in \Omega_M)(\omega \leq_M \omega') \}$

Mae’s Plan A is the set of options that she plans to act on unless they are ruled out by new information. I will refer to $\{ \beta : \langle \beta, \iota \rangle \in \Omega_M \}$ as the set of Mae’s *doxastic alternatives*, $\{ \beta : \langle \beta, \iota \rangle \in \Omega_M^A \}$ as her *top doxastic alternatives*, $\{ \iota : \langle \beta, \iota \rangle \in \Omega_M \}$ as her *practical alternatives*, and $\{ \iota : \langle \beta, \iota \rangle \in \Omega_M^A \}$ as her *top practical alternatives*.

The belief state (B_M) and intention state (I_M) of a model M can be defined as follows:

(37) BELIEF STATES

$B_M = \{ w : (\exists \langle \beta, \iota \rangle \in \Omega_M^A)(w \in \beta) \}$

(38) INTENTION STATES

$I_M = \{ w : (\exists \langle \beta, \iota \rangle \in \Omega_M^A)(w \in \iota) \}$

A model’s belief and intention states are the unions of all of its top-ranked doxastic and practical alternatives, respectively. Notice, also, that it follows from these definitions that the intention state of any cognitive model will be a subset of the belief state.

Beliefs and intentions can be represented as properties of cognitive models—i.e., as functions from models to truth values.

(39) BELIEF IN p

$\lambda M . (\forall \langle \beta, \iota \rangle \in \Omega_M^A)(\beta \subseteq p)$ (EQUIVALENTLY: $\lambda M . B_M \subseteq p$)

(40) INTENTION TO p

$\lambda M . (\forall \langle \beta, \iota \rangle \in \Omega_M^A)(\iota \subseteq p)$ (EQUIVALENTLY: $\lambda M . I_M \subseteq p$)

To believe p is for all of one’s belief worlds to be p worlds—and, equivalently, for all of one’s top doxastic alternatives to entail p . To intend p is for all of one’s intention worlds to be p worlds—and, equivalently, for all of one’s top practical alternatives to entail p .

In keeping with my discussion of partial intentions in Sect. 4, we can say that to partially intend p is for some but not all of one's top practical alternatives to entail p . Formalized as a property of cognitive models:

$$(41) \quad \text{PARTIAL INTENTION TO } P \\ \lambda M . (\exists \langle \beta, \iota \rangle \in \Omega_M^A)(\iota \subseteq p) \text{ and } (\exists \langle \beta, \iota \rangle \in \Omega_M^A)(\iota \not\subseteq p)$$

Any model with this property will also have the property, $\lambda M . I_M \cap p \neq \emptyset$. In other words: if any agent partially intends p , then p is compatible with what they intend. (The converse needn't be true.)

6 Semantics

In Sects. 2–4, I argued that the semantic values of declaratives, strong imperatives, and weak imperatives are beliefs, intentions, and partial intentions, respectively. I will now use what I've already said about cognitive models to formalize these ideas.

I begin with some assumptions about the anatomy of declarative and imperative clauses. First, I will follow one strand of recent work in compositional semantics by taking declarative and imperative clauses to have contents of different semantic types: the content of each declarative is a proposition, while the content of each imperative is a property that is restricted to apply only to the addressee of each context in which it is uttered.²⁵ However, unlike other defenders of this view, I think that there is more than these contents to the semantics of declarative and imperative clauses. In particular, I posit three mood-marker operators: \triangleright for declaratives, $!$ for strong imperatives, and *ixcl*; for weak imperatives, whose semantic values will combine with clausal contents, turning them into properties of cognitive models that represent beliefs, intentions, and partial intentions, respectively.²⁶

²⁵ Versions of this idea have been defended by Portner (2004, 2007); Hausser (1983); Portner et al. (2019); von Stechow and Iatridou (2017); Roberts (2018), and Zanuttini et al. (2012), among others. I adopt it for two reasons. First, it gives us a clear explanation of why imperatives invariably have second-person subjects—albeit subjects that are usually unpronounced in English. (We know they're there, and that they're second person, because imperative subjects can bind second-person pronouns, as in 'give yourself a pat on the back'.) On the view defended by Zanuttini et al. (2012), imperative clauses contain a functional head in their clausal structure called a 'jussive head,' which bears a second-person feature and binds the subject. This leads the clause as a whole to have an addressee-restricted property as its semantic value, and makes imperative subjects invariably second-person (at least in languages like English). Second, this proposal gives us some hope of explaining why imperatives fail to embed in many constructions that embed declaratives (for example, in belief ascriptions). A simple version of this explanation would appeal to the idea that many environments that can compose with proposition-denoting expressions can't embed property-denoting expressions. Portner et al. (2019) defend a more elaborate explanation that appeals to expressive features of imperative subjects that vary cross-linguistically in ways that appear to parallel the cross-linguistic embeddability of imperatives.

²⁶ In these respects, my semantic theory is modeled on that of Charlow (2014), who pioneered the idea of modeling clausal semantic values as properties of a model that represents an agent's beliefs and plans.

(42) DECLARATIVES

$$\llbracket \triangleright \rrbracket = \lambda p . \lambda M . (\forall \langle \beta, \iota \rangle \in \Omega_M^A)(\beta \subseteq p)$$

(43) STRONG IMPERATIVES

$$\llbracket ! \rrbracket^c = \lambda f_{et} . \lambda M . (\forall \langle \beta, \iota \rangle \in \Omega_M^A)(\iota \subseteq \{w : f(a_c) \text{ in } w\})$$

(44) WEAK IMPERATIVES

$$\llbracket iexcl; \rrbracket^c = \lambda f_{et} . \lambda M . (\exists \langle \beta, \iota \rangle \in \Omega_M^A)(\iota \subseteq f(a_c)) \text{ and } (\exists \langle \beta, \iota \rangle \in \Omega_M^A)(\iota \not\subseteq f(a_c))$$

These mood markers play two important roles. First, I have argued that the semantic values of unembedded clauses indicate which sort of mental state a speaker who utters one is trying to instill in their addressee, and this is part of my explanation of why we normally use different clauses to perform different speech acts. Second, we need declaratives and imperatives to share a semantic type when they are conjoined, disjoined, or conditionalized. Mood markers take care of this by raising the contents with which they combine to the same type.

There are two ways that we could think about the grammar of mood markers. First, I could sign on to a tradition that takes every imperative (and possibly also every declarative) to have an operator in its syntax that gives it the special sort of clausal meaning that it has.²⁷ A good reason to avoid this option is the wealth of evidence that declarative clauses typically contribute propositions to the environments in which they embed, such as in attitude ascriptions, under modals, and so on.²⁸ If \triangleright is present in every declarative clause, then we will have to redo the semantics of all of these declarative-embedding expressions. I would like to avoid this conclusion. So, instead of locating mood markers in the syntax of every clause, I will take them to be freely applying operators, which, like type shifters, are grammatically optional. The grammar allows them to be applied to any argument of the right semantic type, but the result may not always be interpretable, whether for semantic or pragmatic reasons. For example, in embedding environments where a declarative clause has to contribute a proposition, \triangleright would render it semantically uninterpretable, and so is, in effect, forbidden. But mixed coordinations will be semantically uninterpretable unless both coordinates have their respective mood markers applied, and so the operators are effectively mandatory in these contexts. And given the pragmatic picture I've painted here, an utterance of an unembedded clause will need the relevant mood marker in position in order to indicate what sort of mental state the speaker intends the addressee to adopt.

Next I need clauses for conjunction, disjunction, and conditionals that allow conjuncts, disjuncts, and consequents to be declarative, imperative, or a mixture of the two. Conjunction can be handled via simple disquotation:

²⁷ See, for example, Starr (2020); Charlow (2014); Kaufmann (2012); Lewis (1970); Davidson (1979); Sadock (1974).

²⁸ Imperatives also embed under attitude verbs in some languages, and probably don't contribute properties of cognitive models when they do. Portner et al. (2019) argue that they typically contribute addressee-directed properties. (But see Kaufmann (2019) for an opposing view on which they contribute modal propositions.)

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

Disjunctions and conditionals are more complicated. The semantic value of a disjunction, for example, mustn't require of a model that it satisfy either disjunct. I argued in Sect. 4 that to satisfy a disjunction is to treat the disjuncts as an exhaustive set of genuine alternatives, which requires not fully satisfying any one disjunct on its own. And to satisfy a conditional is to be in a state such that if one is in that state and also satisfies its antecedent, one is rationally required to satisfy its consequent.

In order to turn these ideas into semantic clauses, it will be useful to think about several kinds of *submodels*. The submodels of a model M are those models that one can get to from M by discarding some of the options that M includes, without changing the order of what remains. Suppose that the model that represents my mental state includes a variety of practical alternatives, each of which entails either that I will write this afternoon or that I will take a nap this afternoon. In that case, my cognitive model is divisible into two submodels, one of which includes the writing alternatives and the other of which includes the napping alternatives. If we assume that neither writing nor napping is my Plan A, then this is the sort of mental state that one might try to put someone in by saying 'Either write or take a nap this afternoon.'

I will annotate submodelhood with a subscripted \downarrow . So, an arbitrary submodel of M will be written M_{\downarrow} . Submodels, in general, are defined as follows:

- (46) SUBMODELS
 A submodel of M is any model M_{\downarrow} such that:

- (i) $\Omega_{M_{\downarrow}} \subseteq \Omega_M$;
- (ii) $\leq_{M_{\downarrow}} \subseteq \leq_M$; and
- (iii) there is no preorder \leq on M_{\downarrow} s.t. $\leq_{M_{\downarrow}} \subset \leq \subseteq \leq_M$.

It will also be useful to single out submodels that keep at least one of the top-ranked options of their supermodels. I will call these 'plan-A-preserving submodels' and annotate an arbitrary such submodel of M as $M_{\downarrow\downarrow}$.

- (47) PLAN-A-PRESERVING SUBMODEL
 M 's plan-A-preserving submodels are the submodels of M in the set,
 $\{M_{\downarrow\downarrow} : \Omega_{M_{\downarrow\downarrow}}^A \cap \Omega_M^A \neq \emptyset\}$

Finally, I define the notion of a *submodel partition*—a set of plan-A-preserving submodels that exclusively and exhaustively carve up a model:

- (48) SUBMODEL PARTITION
 A submodel partition of a model M is a set $\Sigma^{M_{\downarrow\downarrow}} = \{M_{\downarrow\downarrow}^1, M_{\downarrow\downarrow}^2, \dots, M_{\downarrow\downarrow}^n\}$ of plan-

A-preserving submodels of M such that:

- (i) $(\forall M^i, M^j \in \Sigma^{M_\downarrow})(\Omega_{M^i} \cap \Omega_{M^j} = \emptyset)$;
- (ii) $\bigcup(\Omega_{M^1}, \Omega_{M^2} \dots \Omega_{M^n}) = \Omega_M$

With these concepts in hand, we can say that a model satisfies a disjunction just in case it can be submodel-partitioned into a set of submodels, each of which satisfies exactly one of the disjuncts.

(49) DISJUNCTION

$$\llbracket \Phi \text{ or } \Psi \rrbracket = \lambda M . (\exists \Sigma^{M_\downarrow} = \{M^1_\downarrow, M^2_\downarrow\})(\llbracket \Phi \rrbracket(M^1_\downarrow) \text{ and } \llbracket \Psi \rrbracket(M^2_\downarrow))$$

To see how this clause works for imperative disjunctions, consider my example from above. Suppose that you tell me to ‘write or take a nap’. In order to satisfy what you’ve said, I must treat both writing and taking a nap as genuine practical alternatives, without leaving any other alternatives open. This idea is precisified by (49): a model that satisfies your disjunction would be one whose top practical alternatives are all either writing alternatives or napping alternatives, and that includes at least one of each. Now consider a mixed disjunction, such as ‘put back Waverly or I’ll put back Naked Lunch’ (29). To satisfy this disjunction, I would have to be in a state in which all of my top-ranked options constitute either partial intentions to put back Waverly or partial beliefs that the speaker will put back Naked Lunch. As I argued in Sect. 4, this is a state of indecision in which I might find myself when the speaker and I are in the middle of deliberating about how to stay under budget at the book sale—a context in which this sentence would be a natural one to utter.²⁹

To handle conditionals, it will be helpful to pick out just those maximal submodels of a model that satisfy a given sentence. Intuitively, to derive a maximal

²⁹ I have built more than usual into the semantics of disjunction. First, I take it to be a semantic matter that a speaker who utters a disjunction rules out alternatives other than the disjuncts—what Zimmermann (2000, §4.1) calls ‘exhaustivity’. This is also a feature of classical disjunction, but Zimmermann (2000) and Geurts (2005) conclude that it does not belong in the semantics of ‘or’, arguing that it is instead contributed by falling intonation. Second, and this time with Zimmerman and Geurts but contrary to classical disjunction, I have taken it to be a semantic matter that a speaker who utters a disjunction treats both disjuncts as live alternatives—what Zimmermann (2000, §4.1) calls “genuineness.” Several authors have argued that this feature of disjunctive utterances is an implicature rather than part of the meaning of ‘or’ (Grice 1989, 45; Simons 2005, §5.3.1; Alonso-Ovalle 2006, §4.4). I am unconvinced, but I won’t argue the point here. Instead I will point out that it is easy to tweak (49) to remove exhaustivity or genuineness. For example, the following clause removes both:

(49)* DISJUNCTION (WITHOUT GENUINENESS OR EXHAUSTIVITY)

$$\llbracket \Phi \text{ or } \Psi \rrbracket = \lambda M_\downarrow (\llbracket \Phi \rrbracket(M_\downarrow) \vee \llbracket \Psi \rrbracket(M_\downarrow)).$$

My account of Ross’s paradox and imperative free choice both depend on the assumption that ‘or’ has genuineness built into its semantics, and so switching to (49)* would necessitate giving a pragmatic explanation of these phenomena. We could posit a quantity implicature: A speaker who leaves two options open could have closed one off, but chose not to. Since leaving one option out would have been more informative, the speaker must have had a reason not to. The best explanation is that they intended both options to be treated as genuine. (For more sophisticated pragmatic explanations, see Alonso-Ovalle (2005, 2006); Fox (2007); Geurts (2009).)

Φ -satisfying submodel M_{\downarrow}^{Φ} from M is to discard from M as few options as possible such that the result satisfies Φ . (I leave open the possibility that a model may have more than one maximal Φ -satisfying submodel.) Formally:

- (50) MAXIMAL
 Φ -SATISFYING SUBMODEL
 For any sentence Φ , a maximal Φ -satisfying submodel of M is any submodel M_{\downarrow}^{Φ} such that:

- (i) $\llbracket \Phi \rrbracket(M_{\downarrow}^{\Phi})$; and
 (ii) there is no submodel M_{\downarrow} of M such that $M_{\downarrow} \neq M_{\downarrow}^{\Phi}$, M_{\downarrow}^{Φ} is a submodel of M_{\downarrow} , and $\llbracket \Phi \rrbracket(M_{\downarrow})$.

For a model to satisfy a conditional is for all of its maximal antecedent-satisfying submodels to satisfy its consequent:

- (51) CONDITIONAL
 $\llbracket \Phi \rightarrow \Psi \rrbracket = \lambda M . (\forall M_{\downarrow}^{\Phi}) \llbracket \Psi \rrbracket(M_{\downarrow}^{\Phi})$

The idea behind this clause is that you satisfy a conditional just in case each of the minimal ways that you could gain information so as to satisfy the antecedent would rationally commit you to satisfying the consequent. For example, suppose that I say to you, ‘buy me an IPA if they have one’. You will satisfy this sentence if all of your highest-ranked doxastic alternatives in which they have IPAs are paired with practical alternatives in which you buy me one. In other words, my utterance is designed to put you in a state of mind such that learning that they have IPAs (and leaving everything else the same) would rationally compell you to form an intention to buy me one.

7 Consequence and inconsistency

There is an obvious way to define a consequence relation in light of the foregoing:

- (52) CONSEQUENCE
 $\{\Phi_1, \dots, \Phi_m\} \models \Phi_n$ iff $(\forall M : \llbracket \Phi_1 \rrbracket(M) \wedge \dots \wedge \llbracket \Phi_m \rrbracket(M)) (\llbracket \Phi_n \rrbracket(M))$

This says that a conclusion will follow from a set of premises just in case any cognitive model that satisfies all of the premises also satisfies the conclusion. Given that cognitive models are representations of ideally rational minds, (52) formalizes my claim that an argument will strike us as valid when any mind in a state that satisfies all of its premises is rationally required to satisfy its conclusion.

Similarly, the idea that a collection of sentences will strike us as inconsistent if it would be irrational to satisfy all of them at once is embodied by the following definition:

- (53) INCONSISTENCY
 $\text{Inc}\{\Phi_1, \dots, \Phi_n\}$ iff $(\neg \exists M) (\llbracket \Phi_1 \rrbracket(M) \wedge \dots \wedge \llbracket \Phi_n \rrbracket(M))$

A set of sentences is inconsistent just when there is no cognitive model that satisfies all of them at once.

These definitions predict all of the data that I have discussed in earlier sections. In particular, they generate the following predictions (with numbered examples that exemplify each pattern in parentheses to the right):³⁰

- (54) $\{!\phi, \triangleright\phi \rightarrow \triangleright\psi\} \models !\psi$ (4)
- $\{!\phi, \triangleright\phi \rightarrow \triangleright\psi\} \not\models !\psi$ (5)
- $!\phi \not\models !\phi$ or $!\psi$ ROSS'S PARADOX (6)
- $!\phi$ or $!\psi \models \text{imaxcl}; \phi; !\phi$ or $!\psi \models \text{imaxcl}; \psi$ IMPERATIVE FREE CHOICE (7)
- $\text{Inc}\{!\phi, !\neg\phi\}$ (8)
- $\text{Inc}\{\triangleright\neg\phi, !\phi\}$ (9)
- Φ and $\Psi \models \Phi; \Phi$ and $\Psi \models \Psi$ (14)
- $\{\triangleright\phi \rightarrow !\psi, \triangleright\phi\} \models !\psi$ (15)
- $\{\triangleright\phi \rightarrow (!\psi \text{ and } \triangleright\gamma), \triangleright\phi\} \models !\psi$ (16)
- $\{\Phi, \Psi\} \models \Phi \text{ and } \Psi$ (17)
- $\{\triangleright\gamma \rightarrow !\phi, \triangleright\gamma \rightarrow (\triangleright\phi \rightarrow \triangleright\psi)\} \models \triangleright\gamma \rightarrow !\psi$ (18)
- $\text{Inc}\{\triangleright\gamma \rightarrow !\phi, \triangleright\gamma \rightarrow !\neg\phi\}$ (19)
- $\text{Inc}\{\triangleright\gamma \rightarrow \triangleright\neg\phi, \triangleright\gamma \rightarrow !\phi\}$ (20)
- $\text{Con}\{!\phi, \triangleright\neg\phi \rightarrow !\psi\}$ (21)

³⁰ I have represented the disjunctions in both Ross's paradox and imperative free choice as taking wide scope over the mood markers, as in $\uparrow !\phi$ or $!\psi^\uparrow$. A reason to think that this is at least sometimes the right LF is that that mixed imperative–declarative disjunctions pattern with disjunctions of imperatives when it comes to both Ross's paradox and free choice, and there is no single mood marker that can take wide scope in mixed disjunctions. However, VP disjunctions would seem to scope under mood markers, as in 'Post or burn the letter', and so we should probably also block the Ross inference and validate free choice for the LF $\uparrow!(\phi \text{ or } \psi)^\uparrow$. Moreover, we also find versions of Ross's paradox and free choice that involve indefinites, which presumably scope under mood (e.g. Ross: 'Help yourself to that cookie on the left' $\not\models$ 'Help yourself to a cookie').

Here is one way to extend my semantics to make these predictions. First, we model disjunction and indefinites as alternative-introducing expressions, and we take sentence radicals to denote sets of alternatives rather than single propositions or properties (Alonso-Ovalle 2006; Simons 2005; Aloni 2003, 2007; Charlow 2019). For example, the content of 'have coffee or tea' would be a set that includes the property of having coffee and the property of having tea. Second, we adjust our semantics for mood markers so that they take alternative sets rather than propositions or properties as inputs. For example, here is a revised clause for !:

$$(43)^* \quad \text{STRONG IMPERATIVE ALTERNATIVE SEMANTICS } \llbracket ! \rrbracket^c = \lambda P . \lambda M . (\forall f \in P)(\exists \langle \beta, \iota \rangle \in \Omega_M^A)(\iota \subseteq f(a_c)) \wedge (\forall \langle \beta, \iota \rangle \in \Omega_M^A)(\exists f \in P)(\iota \subseteq f(a_c))$$

This clause would make $!(\phi \text{ or } \psi)$ equivalent to $!\phi$ or $!\psi$, and could generalize to indefinites as well as disjunction. Of course, this is just a rough sketch of an idea that would need much more work.

8 Conclusions and Loose Ends

We are agents who make complex plans, who aspire to rational coherence in our thought and action, and who are at least sometimes sensitive to our failures to meet this standard. These are deep features of human cognition, but they can be found on the surface of our language use, having bubbled up through the computational mechanisms that we use to encode and decode the meanings of sentences. We speak in order to influence each others' states of mind, and clauses of different types are used to instill different kinds of mental states—declaratives for beliefs, imperatives for intentions, and complex multi-clausal constructions for complex states that sometimes arise in theoretical and practical reasoning. These uses are encoded into sentences' meanings, and this allows us to construct natural-language arguments that track the rational relations that hold among the mental states that sentences are used to produce.

Or so I have argued. Should you be convinced? You might want to wait for a data-driven blow-by-blow comparison of my formal semantics with the various alternatives currently alive in the literature—something that would have at least doubled the length of this paper, and that I therefore left out in favor of painting as detailed a portrait of my own view as I could manage. Of course, I *have* emphasized what I take to be my theory's main advantages. It explains the striking parallels between the rules governing imperative inference and the principles governing practical rationality. It explains the awkwardness of using an imperative to tell someone to do something over which they have no control. It makes sense of the surprising ways that imperatives and declaratives can mix, and also explains the severe restrictions on the felicitous use of these mixtures. It explains why imperative arguments deviate from classical logic. Most importantly (at least in my view), my account explains all of these things in terms of independently motivated underlying features of practical rationality—something to which few previous semantic theories have aspired. Still, much remains to do.

References

- Aloni, M. (2003). Free choice in modal contexts. In M. Weisgerber (Ed.), *Proceedings of Sinn und Bedeutung 7*, chapter 2, pp. 25–37. University of Konstanz.
- Aloni, M. (2007). Free choice, modals, and imperatives. *Natural Language Semantics*, 15, 65–94.
- Alonso-Ovalle, L. (2005). Distributing the disjuncts over the modal space. In L. Bateman, C. Ussery (Ed.), *Proceedings of the 35th North East Linguistics Society Conference*, pp 75–86. GSLA, Amherst: University of Massachusetts.
- Alonso-Ovalle, L. (2006). *Disjunction in Alternative Semantics*. Amherst: University of Massachusetts (PhD thesis).
- Anscombe, G. E. M. (1963). *Intention* (2nd ed.). Oxford: Blackwell.
- Audi, R. (1973). Intending. *Journal of Philosophy*, 70, 387–403.
- Barker, C. (2012). Imperatives denote actions. In A. A. Guevara, A. Chernilovskaya, R. Nouwen (Ed.), *Proceedings of Sinn und Bedeutung 16*, chapter 5. MIT Working Papers in Linguistics.
- Bratman, M. (1987). *Intention, plans, and practical reason*. Cambridge: Harvard University Press.
- Bratman, M. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 327–341.

- Bratman, M. (1999). *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- Bratman, M. (2007). *Structures of Agency: Essays*. Oxford: Oxford University Press.
- Bratman, M. (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Broome, J. (1999). Normative requirements. *Ratio*, *XI*, *I*(4), 398–419.
- Broome, J. (2013). *Rationality Through Reasoning*. West Sussex: Wiley-Blackwell.
- Charlow, N. (2014). Logic and semantics for imperatives. *Journal of Philosophical Logic*, *43*, 617–664.
- Charlow, N. (2018). Clause-type, force, and normative judgment in the semantics of imperatives. In D. Fogal, D. Harris, & M. Moss (Eds.), *New Work on Speech Acts*. Oxford: Oxford University Press.
- Charlow, S. (2019). The scope of alternatives: Indefiniteness and islands. *Linguistics and Philosophy*, *43*(4), 427–472.
- Chisholm, R. (1963). Contrary-to-duty imperatives and deontic logic. *Analysis*, *24*(2), 33–36.
- Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, *42*, 213–261.
- Condoravdi, C., & Lauer, S. (2012). Imperatives: Meaning and illocutionary force. In C. Pinõn, (Ed.), *Empirical Issues in Syntax and Semantics 9: Papers From the Colloque de Syntaxe et Sémantique à Paris 2011*, pp 37–58.
- Culicover, P. W., & Jackendoff, R. (1997). Semantic subordination despite syntactic coordination. *Linguistic Inquiry*, *28*(2), 195–217.
- Davidson, D. (1979). Moods and performances. In A. Margalit (Ed.), *Meaning and Use* (pp. 9–20). Dordrecht: D. Reidel.
- Ferrero, L. (2009). Conditional intentions. *Noûs*, *43*(4), 700–741.
- von Fintel, K., & Iatridou, S. (2017). A modest proposal for the meaning of imperatives. In A. Arregui, M. Rivero, & A. P. Salanova (Eds.), *Modality Across Semantic Categories, chapter 13* (pp. 288–319). Oxford: Oxford University Press.
- Fogal, D. (2020). Rational requirements and the primacy of pressure. *Mind*. (forthcoming).
- Fox, D. (2007). Free choice disjunction and the theory of scalar implicatures. In U. Sauerland & P. Stateva (Eds.), *Presupposition and implicature in compositional semantics* (pp. 71–120). New York: Palgrave Macmillan.
- Geurts, B. (2005). Entertaining alternatives: Disjunctions as modals. *Natural Language Semantics*, *13*, 383–410.
- Geurts, B. (2009). Scalar implicature and local pragmatics. *Mind and Language*, *24*, 51–79.
- Grice, H. P. (1968). Utterers meaning, sentence-meaning, and word-meaning. *Foundations of Language*, *4*(3), 225–242.
- Grice, H. P. (1969). Utterers meaning and intention. *The Philosophical Review*, *78*(2), 147–177.
- Grice, H. P. (1971). Intention and uncertainty. *Proceedings of the British Academy*, *57*, 263–279.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Harman, G. (1976). Practical reasoning. *The Review of Metaphysics*, *29*, 431–463.
- Harman, G. (1986). *Change in View: Principles of Reasoning*. Cambridge: MIT Press.
- Hausser, R. (1980). Surface compositionality and the semantics of mood. In J. Searle, F. Kiefer, & M. Bierwisch (Eds.), *Speech Act Theory and Pragmatics* (pp. 71–96). Dordrecht: Reidel.
- Hausser, R. (1983). The syntax and semantics of English mood. In F. Kiefer (Ed.), *Questions and Answers* (pp. 97–158). Dordrecht: Reidel.
- Hill, T. E. (1973). The hypothetical imperative. *The Philosophical Review*, *82*(4), 429–450.
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. New York: Cornell University Press.
- Holton, R. (2008). Partial belief, partial intention. *Mind*, *117*(465), 27–58.
- Holton, R. (2011). *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
- Kaufmann, S., & Schwager, M. (2009). A unified analysis of conditional imperatives. In Cormany, E., Ito, S., Lutz, D., (Ed.), *Proceedings of SALT 19*, pp 239–256. Columbus: Ohio State University.
- Kaufmann, M. (2012). *Interpreting Imperatives*. Dordrecht: Springer.
- Kaufmann, M. (2019). Fine-tuning natural language imperatives. *Journal of Logic and Computation*, *29*(3), 321–348.
- Keshet, E. (2012). Focus on conditional conjunction. *Journal of Semantics*, *30*(2), 211–256.
- Keshet, E., & Medeiros, D. J. (2019). Imperatives under coordination. *Natural Language and Linguistic Theory*, *37*, 869–914.

- Klinedinst, N., & Rothschild, D. (2012). Connectives without truth tables. *Natural Language Semantics*, 20, 137–175.
- 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.
- Lewis, D. K. (1970). General semantics. *Synthese*, 22(1/2), 18–67.
- Neale, S. (2016). Silent reference. In G. Ostertag (Ed.), *Meanings and Other Things: Essays in Honor of Stephen Schiffer* (pp. 229–342). Oxford: Oxford University Press.
- Portner, P. (2004). The semantics of imperatives within a theory of clause-types. In Watanabe, K., Young, R., (Ed.), *Proceedings of SALT 14*. Washington: CLC Publications.
- Portner, P., Pak, M., & Zanuttini, R. (2019). The speaker-addressee relation in imperatives. In Bondarenko, T., Davis, C., Colley, J., Privoznov, D., (Ed.), *MIT Working Papers in Linguistics 90: Proceedings of the 14th Workshop on Altaic Formal Linguistics (WAFL14)*. MIT Working Papers in Linguistics, Cambridge.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15, 351–383.
- Portner, P. (2012). Permission and choice. In G. Grewendorf & T. Zimmermann (Eds.), *Discourse and Grammar: From Sentence Types to Lexical Categories* (pp. 43–68). Berlin: De Gruyter Mouton.
- Roberts, C. (2012). Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5, 1–69.
- Roberts, C. (2018). Speech acts in discourse context. In D. Fogal, D. Harris, & M. Moss (Eds.), *New Work on Speech Acts*. Oxford: Oxford University Press.
- Ross, A. (1941). Imperatives and logic. *Theoria*, 25(7), 53–71.
- Sadock, J. (1974). *Toward a Linguistic Theory of Speech Acts*. New York: Academic Press.
- Schroeder, M. (2004). The scope of instrumental reason. *Philosophical Perspectives*, 18, 337–364.
- Searle, J. (1969). *Speech Acts*. London: Cambridge University Press.
- Simons, M. (2000). *Issues in the Semantics and Pragmatics of Disjunction*. New York: Garland.
- Simons, M. (2005). Dividing things up: The semantics of or and the modal/or interaction. *Natural Language Semantics*, 13, 271–316.
- Starr, W. B. (2010). *Conditionals, Meaning and Mood*. PhD thesis, Rutgers University.
- Starr, W. (2020). A preference semantics for imperatives. *Semantics and Pragmatics*, Early Access
- Starr, W. B. (2014). Mood, force, and truth. *Protosociology*, 31, 160–180.
- Stokke, A. (2014). Insincerity. *Nouÿs*, 48(3), 496–520.
- 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.
- Zanuttini, R., Pak, M., & Portner, P. (2012). A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects. *Natural Language and Linguistic Theory*, 30(4), 1231–1274.
- Zimmermann, T. E. (2000). Free choice disjunction and epistemic possibility. *Natural Language Semantics*, 8, 255–290.