

Conditional preferences and practical conditionals

Nate Charlow

© Springer Science+Business Media Dordrecht 2014

Abstract I argue that taking the Practical Conditionals Thesis (PCT) seriously demands a new understanding of the semantics of such conditionals.

Practical Conditionals Thesis: A practical conditional [*if A*][*ought(B)*] expresses *B*'s conditional preferability given *A*

Paul Weirich has argued that the conditional utility of a state of affairs *B* on *A* is to be identified as the degree to which it is desired under indicative supposition that *A*. Similarly, exploiting the PCT, I will argue that the proper analysis of indicative practical (as well as imperative) conditionals is in terms of what is planned, desired, or preferred, given suppositional changes to an agent's information. Implementing such a conception of conditional preference in a semantic analysis of indicative practical conditionals turns out to be incompatible with any approach which treats the indicative conditional as expressing non-vacuous universal quantification over some domain of relevant antecedent-possibilities. Such analyses, I argue, encode a fundamental misunderstanding of what it is to be best, given some condition. The analysis that does the best vis-à-vis the PCT is, instead, one that blends (i) a Context-Shifty account of indicative antecedents with (ii) an Expressivistic, or non-propositional, treatment of their practical consequents.

Keywords Indicative conditionals · Practical conditionals · Anankastic conditionals · Conditional imperatives · Conditional preference · Prioritizing modality · Dynamic semantics · Expressivism

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

1 Introduction

The degree to which a rational agent wants a state of affairs A can depend on what is the case about A , in at least two ways. First, if A causally conduces to some distinct thing she wants. The degree to which an agent wants, say, to quit smoking can depend on whether quitting would make her live longer. Second, if A is constitutive of something desirable. The degree to which an agent wants, say, to be virtuous can depend on whether virtue is an essential part of a good life. Or, more mundanely, the degree to which an agent wants something can depend on whether it is something she wants at all. The same goes for *non-degreed* preferences. Whether something is preferred simpliciter can depend, in both of the above ways, on what is the case about A .

Decision theorists have plausibly argued that such dependencies demand introduction of a notion of *conditional utility* into decision-theoretic frameworks. By exploiting what I term the “Practical Conditionals Thesis,” this essay argues that the notion of conditional utility—more precisely, the related notion of conditional preferability—is important, not just for decision theorists interested in a realistic and comprehensive decision theory—something on which I spend no time here—but also for anyone interested in the meaning of natural language constructions that function to *express features of conditional preferences*, including practical conditionals (*if A, you ought to see to it that B*¹) and related conditional imperatives (*if A, see to it that B*).

Practical Conditionals Thesis (PCT)

A practical conditional [*if A*][*O(B)*] expresses B 's conditional preferability given A

Weirich (1980) has argued compellingly that the conditional utility of a state of affairs B on another state of affairs A is to be identified as the degree to which it is desired under *indicative supposition* that A . Similarly, exploiting PCT, I argue that the proper analysis of indicative practical and imperative conditionals is in terms of what is planned, desired, or preferred, given monotonic changes to an agent's information (loosely, suppositions). Though my main focus in this paper is on questions of linguistic meaning, also interesting, and susceptible of the sort of methodology that I employ, is the related question of how to represent the conventional *cognitive force* of constructions that express conditional preference. Here, too, Weirich's idea bears fruit.

Implementing Weirich's conception of conditional preference in a semantics for practical conditionals turns out to be incompatible with many popular accounts of their meaning. These include Wide-Scope accounts, the philosophically standard “Variably Strict” quantificational possible worlds semantics of Stalnaker and Lewis, and even sophisticated versions of Context-Shifting accounts (on which indicative antecedents function to shift the context of interpretation for material in their consequents). In fact, I will ultimately argue that it is *prima facie* incompatible with any *quantificational* semantics for indicative conditionals—any approach which treats the indicative

¹ For simplicity, I will generally obscure the semantic distinction between weak practical necessity modals (*should, ought*) and strong practical necessity modals (*must, have to*). The sentential operator O functions to represent a generic prioritizing necessity modality.

conditional as expressing non-vacuous universal quantification over some domain of relevant antecedent-possibilities.² The analysis that does the best job with simple practical conditionals like *if you want A, you should see to it that A* is an account that blends (i) a Context-Shifty account of indicative antecedents with (ii) an Expressivistic treatment of their practical consequents. To put the point bluntly, if we treat practical conditionals as expressing conditional preferability, competitor analyses of such indicative conditionals encode a fundamental misunderstanding of what it is to be *best, given some condition*.

I begin by introducing the notion of conditional utility and briefly rehearsing Weirich's arguments for treating conditional utilities as suppositional degrees of desire. Next I profile the class of conditionals that, I claim, can be fruitfully theorized about in terms of conditional preferability. The class includes conditionals commonly referred to as "anankastic", but other practical conditionals besides these. I then describe three pieces of data that any account of practical conditionals must accommodate. First, modalities expressing practical necessity in the relevant readings of the relevant conditionals should be given narrow scope. Second, practical conditionals do not license a phenomenon I refer to as Bootstrapping—inferring from *if you want A, you should see to it that A* that you should see to it that A, when you simply happen to want A. Third, the antecedents of practical conditionals are non-monotonic in a distinctive way—the inference from *if you want A, you should see to it that C* to *if you want A and B, you should see to it that C* is illegitimate, for reasons distinctive to practical conditionals.

The subsequent discussion introduces three ways of giving the relevant modalities narrow scope, or Narrow-Scoping—the Variably Strict approach, Context-Shifty Cognitivism, and Context-Shifty Expressivism—and argues that all manage to avoid the sorts of concerns that trouble wide-scope accounts. However, I go on to argue that Variable Strictness and Context-Shifty Cognitivism—both of which treat practical conditionals as expressing, roughly, *known* preferability, on update with the relevant condition—validate either bootstrapping or downward monotonicity for practical antecedents. Context-Shifty Expressivism—which treats practical conditionals as expressing preferability, on update with the relevant condition—validates neither. Various damage-control strategies are considered for Variable Strictness and Context-Shifty Cognitivism. I argue that none are convincing.

In the final section, I consider how an Expressivist account of the *cognitive force* of practical conditionals might look. I suggest that entertaining a practical conditional is, for the Expressivist, a matter of *consulting one's preferences*, on update with its antecedent. This, importantly, is distinct from consulting one's *beliefs about* one's preferences, on update with the antecedent. Correspondingly, deciding that a practical conditional is the case is a matter of ending up in a cognitive state characterized by a specific, suppositional preference. This, I argue, is distinct both from (i) deciding *that* there is such a suppositional preference (hence distinct from a Cognitivist account of endorsement of practical conditionals), and (ii) *coming to have* such a suppositional preference. This is a good thing. Any account that conflates a decision to accept a

² Gillies (2010: Sect. 4) shows why any plausible quantificational analysis of the indicative conditional will need to treat it as expressing this sort of quantification.

practical conditional with either (i) or (ii) is incorrect as an account of their cognitive force; the former because it misconstrues the subject matter of such a decision, the latter because it makes it mysterious how a decision to accept a practical conditional can *resolve a question* for the agent (namely, the issue of *whether* the practical conditional is the case). On our favored account of the cognitive force of practical conditionals, it is possible to entertain whether a practical conditional is the case (hence, accepting a practical conditional is potentially issue-resolving for an agent). This discussion speaks to a few largely unaddressed, but central, issues in the literature on expressivism. Philosophers of language and meta-ethicists who are interested in expressivism, but who find the argumentation of the penultimate section difficult, will still, I think, find value in it.

2 Conditional utility

Weirich (1980) considers three ways of conceptualizing the notion of conditional utility.³ In this section, I describe them briefly, and lay out Weirich's arguments in favor of treating conditional utility as degree of desire under indicative supposition. Since conditional utility, per se, is not our focus, the discussion is spare (but still, I hope, suggestive).

Proposal 1: Reduction. According to the proposal that we will call *reduction*, conditional utility is identified with the utility of the corresponding conditional: $u(A|B) = u(B \rightarrow A)$.⁴ (Here and throughout, u is a relative utility function, i.e., if $u(X)$ is positive, then X represents an improvement on the status quo.) Weirich objects to reduction on the following grounds (1980, p. 704).

Cookie

John wants a cookie (c), and also wants milk (m) if he has a cookie. John has decided not to have any milk, since he doesn't know whether or not he will get a cookie, and he only likes milk when he also has a cookie to eat with it.

Letting u be John's relative utility function, $u(c)$ and $u(m|c)$ are each positive. However, $u(c \rightarrow m)$ is negative; given what John knows, the only way that $(c \rightarrow m)$ is if $\neg c$, and John *really* wants a cookie. Nevertheless, conditional on getting a cookie, he *really* wants some milk, so $u(m|c)$ is still positive. So, in general,

$$u(A|B) \neq u(B \rightarrow A)$$

Proposal 2: Conditional Prediction. On the second proposal Weirich entertains, an individual's conditional utility for A on B is identified with the utility that A would have, were the individual to learn that B . In particular, if u_B gives the agent's

³ For other discussions, see Jeffrey (1983), Bradley (1999) and Joyce (2000).

⁴ Reduction is a common move when attempting to understand conditional states of mind. Given what we know about conditional states of mind, it is hard to understand why. It is well-known that conditional probability, e.g., does not reduce to the unconditional probability of a proposition (Lewis 1976). There is no reason, save lack of imagination, to suppose that bearing \mathcal{A} toward X conditional on Y should ever be understood as bearing A toward some proposition Z .

utility function, *were she to learn B*, then $u(A|B) = u_B(A)$. This proposal involves no reduction of conditional utility to regular utility. The conditional utility an agent assigns to *A* conditional on *B* depends on the properties of a utility function that the agent would have—possibly (indeed, typically) distinct from her actual utility function—were she to conditionalize on *B*.

Weirich (1980) objects to this proposal on the grounds that “conditional utilities may differ from the corresponding conditionally predicted utilities because the relative utility of an action given a condition may depend on one’s being ignorant of the condition” (p. 705). Here is an adaptation of a case he gives to illustrate the point.

Case 1.

John wants to be surprised. He would be surprised to observe flowers in bloom in Toronto in March. If there are flowers in bloom in Toronto in March (*f*) and John were to go to Toronto then (*t*), John would observe that *f*. However, if John were simply to learn *f*, he would, of course, *not* be surprised to observe *f*. And going to Toronto is, by itself, kind of a hassle to John.

Two things are obvious from the description of the case. First, *t* is desirable conditional on *f*, i.e., $u(t|f)$ is positive. Second, $u_f(t)$ —the relative utility of going to Toronto once John conditionalizes on *f*—is negative. So, in general,

$$u(A|B) \neq u_B(A)$$

Digression: The Ramsey Test. A quick digression, by way of (i) establishing a preliminary link between conditional utility and practical conditionals, (ii) motivating the next proposal about conditional utility. Something like the above point—that an agent cannot, in general, evaluate conditional entities (propositions, utilities, etc.) by evaluating what would hold were she to conditionalize on the relevant condition—is familiar from various discussions of the Ramsey Test. The naive formulation of the Ramsey Test says: to evaluate an indicative conditional of the form [*if p*][*q*] pretend you know *p*, then assess whether *q*. The conditional [*if p*][*q*] is acceptable just if *q* is accepted when one pretends that she knows *p*. This formulation runs into well-known problems. For instance, (1) is often acceptable.

(1) If my spouse is cheating, I won’t find out

But when one evaluates whether one will find out about their spouse cheating when they pretend they know this, one finds the only answer is ‘yes’. So the naive formulation of the Ramsey Test cannot give a generally appropriate procedure for evaluating an indicative conditional. This is structurally identical to the problem for Proposal 2 described above.

A better version of the Ramsey Test⁵ says that for an indicative [*if p*][*q*]: (i) the antecedent *characterizes a property of an information state*: the property of accepting *p*, (ii) to evaluate the indicative, move to the nearest information state that satisfies

⁵ One popularized by Irene Heim and a family of dynamic semantic treatments of conditionals. See, e.g., Karttunen (1973), Heim (1983), Yalcin (2007) and Gillies (2010).

that property (which we can gloss roughly as *supposing that p*), and then evaluate whether *q*, from the vantage point of that enriched information state. This is crucially different from evaluating, were one to learn that *p*, whether *q*, or evaluating, if you pretend you know that *p*, whether *q*. In supposing that *p*, one leaves it open whether or not one *also* accepts or knows *p*, under this supposition. If one entertains *p as a supposition*, one is not thereby committed to entertaining the propositions that one has learned *p*, one knows *p*, or one accepts *p*, as holding under this supposition. In short: one can perfectly well suppose that: *p* and I do not know (accept) that *p*.

Proposal 3: Suppositional Degree of Desire. All this seems to go, *mutatis mutandis*, for conditional utilities. Proposal 2 errs by suggesting that determining something's relative utility, conditional on *p*, is a matter of evaluating its relative utility on the supposition that *p* is known or learned. This would seem to suggest an account which identifies something's conditional utility as its unconditional utility *under the supposition that the relevant condition holds* (where evaluating under the supposition that *p* is understood in terms of simply having an information state that accepts *p*, but which may fail to accept that it accepts that *p*). And, indeed, that's exactly what Weirich goes on to suggest. In particular, if $u_{|s|}$ gives the agent's "hypothetical" utility function, when she supposes *s*, then

$$u(A|B) = u_{|B|}(A)$$

Pretty clearly Proposal 3 gets the right results about the cases that troubled Proposals 1–2. You can still want a cookie and want milk, conditional on getting a cookie, even when you've decided that you do not want to have milk. Proposal 3 explains this by noting that your preferences may change under indicative supposition. Supposing you have a cookie, of course you want milk to go with it. Additionally, you can want to be surprised and want to go to Toronto, under the supposition that the flowers are in bloom in Toronto in March, because, under this supposition, if you go to Toronto then, you will be surprised.

Notice that the unviability of Proposal 1 means that the notion *degree of desire under supposition that the condition holds* cannot be understood as the degree of desire that the corresponding material implication hold. Rather, conditional utilities are understood, irreducibly, as degrees of desire, read off from the agent's preferences under the relevant supposition.⁶

Signpost. So much for conditional utility, as such. My purpose in the remainder of this paper is to explore how a proper understanding of conditional utility can impact our understanding of the semantics of constructions in natural language that are used to *express conditional preferences*. Interestingly, as we will see, issues like those sketched in this section arise for such constructions. This is of fundamental significance for the semantic analysis of such conditionals.

⁶ Precising this proposal gives rise to complications, but addressing them would bog us down in decision-theoretic issues that are not strictly relevant to my project here. The proposal that we eventually articulate for practical conditionals will be sufficiently precise and well-developed to constitute a constructive account of their meaning and cognitive force.

3 Profiling practical conditionals

Phenomena similar to those presented above arise in the case of practical conditionals. In this section, we will (i) describe the central characteristic that such conditionals share, (ii) see why the analysis of conditional utility might be important for the semantics of such conditionals. While the claims I make here are liable to be controversial, the point of this section is to *contextualize* the paper's overall focus and argument. As I will explain, it can perform this function even for readers who do not find the central claims compelling.

Practical Conditionals, Neutrally. The main target of this paper is conditionals like:

- (2) If you want to go to Harlem, you should take the A-train
- (3) If you want to have a family, you should be married
- (4) If you want to be a professor, you should be a professor
- (5) If that cheese tastes good, you should eat it
- (6) If the third rail is live, you should avoid it

Though these conditionals exhibit important differences (more on these shortly), they share important characteristics. Preliminarily: each has an antecedent that is, broadly speaking, relevant to what the agent in question should do. Each has a consequent containing a prioritizing necessity modal (a modal that, in contrast with, e.g., epistemic necessity modals, is, in more or less direct fashion, supposed to have *motivational*, rather than cognitive, *force for the agent that is its subject*; cf. Schwager (2006), Portner (2007), Kaufmann and Schwager (2009)). Finally, the modal's interpretation is sensitive, in a sense to be precisified, to the information borne by the antecedent. Such conditionals, I will suggest, form part of a semantic natural kind that I will designate with the label *practical conditionals*.

Conditional Preferability. What unites (2)–(6)? Why regard them as a semantic natural kind? An intuitive suggestion: all of these conditionals express that the relative desirability of the modal's object, conditional on the information expressed by the antecedent, is high. This claim about the meaning of such conditionals follows from a more general claim about practical conditionals, what I will term the Practical Conditionals Thesis.

Practical Conditionals Thesis (PCT)

A practical conditional $[if A][O(B)]$ expresses B 's conditional preferability given A

Four short notes about this. First, as I will explain more below, I do not mean to claim that every conditional of the form $[if A][O(B)]$ is practical. Indeed I deny this. Practical conditionals are not individuated by superficial form, rather by meaning. Second, PCT may be understood to function as a stipulative definition of the kind *practical conditionals*; the degree to which it is interesting will then depend on how

large the class of conditionals whose meaning can be fruitfully understood in the way it suggests is. (More about what this class looks like shortly.) Third, preferability is to be understood as a comparative notion: B is preferable just if, for any relevant alternative C to B , B is preferred to C . B is preferable conditional on A just if, for any C such that C is a relevant alternative to B if A , B is conditionally preferred (on A) to C . Fourth, preferability in the relevant sense is generally determined by the grammatical subject of the prioritizing modal in the consequent—generally, but not necessarily, the speaker’s addressee. So, for example, the sense of preferability relevant for the interpretation of (2) is preferability *for* the speaker’s addressee.⁷

It is perhaps instructive to compare PCT with Adams’ Thesis (Adams 1975), a plausible understanding of which says that a non-practical indicative conditional expresses that the probability of the consequent, conditional on the antecedent, is high (although, for some qualifications, see Hájek 2012). While Adams arguably identifies the probability of an entire conditional with a conditional probability, the PCT makes no claim about the utility of the entire conditional. But the intuitions that support the idea that a non-practical indicative conditional expresses that the probability of the consequent, conditional on the antecedent, is high also, I think, support the idea that a practical indicative conditional expresses that the desirability of the consequent, conditional on the antecedent, is high. In general, indicative antecedents serve to conditionalize the state of mind conventionally expressed by their consequents; since an atomic declarative q conventionally expresses, e.g., a high degree of confidence in the proposition that q , the indicative $[if\ p][q]$ expresses a high degree of confidence in q conditional on p . Similarly, since a prioritizing necessity modal $O(q)$ expresses q ’s comparative preferability, the indicative $[if\ p][O(q)]$ expresses q ’s comparative preferability, conditional on p .

In this paper, I will not be relying on PCT’s truth, so much as trying to deploy PCT as a theoretical tool with the following dual function:

- **Diagnostic:** a device for determining why certain analyses of the semantics of practical conditionals seem to generate incorrect meanings for practical conditionals
- **Heuristic:** a device for guiding the development of an analysis of the semantics of practical conditionals that does better

While I take PCT’s suitability for this function, when paired with a Weirich-inspired understanding of conditional preferability, to be a major point in favor of each (and the thesis that practical conditionals constitute a semantic natural kind), the semantics I develop for practical conditionals is neutral on this. Insofar as the reader agrees with the judgments I use PCT to reach (and the analysis these judgments ultimately support), PCT’s truth is immaterial.

The Scope of PCT. I’ve given only a vague general description of the class of conditionals whose meaning is characterized by PCT: conditionals of the form $[if\ A][O(B)]$

⁷ It is important to note that this should not be taken to imply that the relevant modal has a *teleological* meaning (at least in the standard sense of that notion, on which the claim that x should Y has a teleological meaning if it expresses that Y is best given x ’s goals and desires). What is preferable-for- x -given- A can be partly constituted by the goals and desires that x *ought to have*, given A , as we will see in Sect. 4.3.

meeting certain other (also vague) constraints (relevance of the antecedent to the desirability of the prejacent of a prioritizing necessity modal in the consequent). Can I be less vague?

Before I try, a word about the ambitions of this paper. My primary aim, as I have hinted, is to use PCT to think about how to semantically represent salient readings of certain kinds of conditionals—conditionals that have an intuitively practical meaning (as defined by PCT). In realizing this aim, I do not require a general description of the class *practical conditionals*. I require only a well-defined data set, a description of its central properties, and a semantics which accounts for those properties. For this purpose, it will suffice to understand the class of practical conditionals minimally, as just those conditionals for which my semantics is the correct semantics.

Of course, in assessing the scope and importance of PCT for the study of conditionals of the form $[if\ A][O(B)]$ generally, readers may be interested in how to tell whether any such conditional is practical, besides consulting their (perhaps less-than-perfectly-reliable) intuitions about whether a conditional expresses conditional preferability. In Sect. 4, I will give three defining (if not definitive) characteristics of practical conditionals (or, more precisely, practical readings of conditionals of the form $[if\ A][O(B)]$) and argue that these are attributable to the fact that the conditionals in question express conditional preferability. The characteristics, roughly, are these: (i) they resist representation with wide-scoped modal operators, (ii) they allow the inference⁸ of their consequents just when the information expressed by their antecedents is available to or accepted by the agent that is the subject of their prioritizing modal, and (iii) their antecedents function, in part, to hypothetically modify the preferences or criteria against which outcomes are ranked according to desirability or preferability.

Anankastics. In the meantime, let me say something about what practical conditionals, as a class, are *not*. Recent work has largely taken the semantically interesting class in this neighborhood to be conditionals with an “anankastic” meaning—anankastic conditionals (see esp. von Fintel and Iatridou 2005).⁹ I will partly follow this trend in this paper: anankastic conditionals represent my main source of data. Nevertheless, it would be a mistake to think that PCT is a claim *about* anankastic conditionals (or that the semantic natural kind it carves out is the same one carved out by anankastics): many anankastic conditionals clearly cannot be profitably treated under the rubric of PCT (and do not exhibit defining characteristics of practical conditionals), while many non-anankastic conditionals clearly can (and do exhibit defining characteristics of practical conditionals).

⁸ Here and throughout the paper, I am willfully (but harmlessly) sloppy about the distinction between legitimate inferences (a matter, strictly speaking, of proper concern for normative epistemology, rather than semantics) and semantic entailments.

⁹ The label, due originally to von Wright (1963), is misleading. Per Liddell and Scott’s standard *Greek-English Lexicon*, *anankastic* means, roughly, concerning force, constraint, or compulsion (thus the WHO’s designation of obsessive-compulsive disorder as “anankastic personality disorder”; see <http://www.mentalhealth.com/icd/p22-pe10.html>). Weak prioritizing necessity modals (*should* or *ought* as opposed to *must* or *have to*) generally express that their subject is *not* constrained or compelled to perform the referenced action (cf. von Fintel and Iatridou 2008).

Anankastic conditionals are generally individuated by their paraphrasability with (or equivalence to) a purpose *to*-clause: they are composed of a subordinate *if*-clause that expresses a purpose (or goal, end, *telos*, etc.), and a teleologically modal matrix clause saying what ought or has to obtain to realize that purpose (cf. Sæbø 2001; von Fintel and Iatridou 2005).¹⁰ Thus the felt equivalence of (2) and:

(7) To go to Harlem, you should take the A-train

According to the paraphrasability criterion, however, (3) dubiously qualifies as anankastic, while (4) does not.

(8) ??To have a family, you should be married

(9) #To be a professor, you should be a professor

In addition, since the antecedents of (5) and (6) do not express goals, hence are not paraphrasable with a purpose *to*-clause, they are clearly not to be typed as anankastic. Since there is an important semantic commonality running through (2)–(6), restricting our theoretical focus to conditionals classified as anankastic according to the paraphrasability criterion risks missing an important semantic generalization.¹¹ Even more generally, (2)–(6) comprise a semantic natural kind that includes, not only conditionals whose antecedents express *goals*, but additionally conditionals whose antecedents express *circumstances* relevant to the desirability of the state of affairs enjoined by the consequent. (Although most of the data I discuss involves practical conditionals with goal-expressing antecedents, readers who make it to the end of the paper should satisfy themselves that the analysis I eventually develop yields a *prima facie* tractable analysis of practical conditionals with non-goal-expressing antecedents.)

Arguably, then, not every practical conditional is anankastic. Is every anankastic conditional (so-classified according to the paraphrasability criterion) a practical conditional? No. Anankastic conditionals can express a pure (what we might term) relationship of necessity between states of affairs—i.e., a state of affairs in which one state of affairs is simply a (metaphysically, logically, conceptually, physically,

¹⁰ Anankastic conditionals are sometimes also individuated according to rough semantic criteria—e.g., fulfilling the consequent is the best or only *means to* the fulfillment of the goal expressed by the antecedent. By this criterion, only (2) is anankastic. Being married is not a means to having a family; it is (partly) constitutive of it. (To get a handle on the distinction: lifting weights is a means to having a better physique, but is not a means to doing physical activity; rather, it is partly constitutive of it.) Being a professor is not a means to being a professor; it is (wholly) constitutive of it. Again, restricting our focus to anankastics, when anankastics are individuated in this way, misses the important semantic commonality running through (2)–(6).

¹¹ Thus the literature's special focus on anankastics, and implicit assumption that their analysis requires a semantics that may not work for practical conditionals in general, is arguably distorting. Specifically, work on the semantics of anankastics has focused on a phenomenon that, while important, does not figure centrally in the treatment of practical conditionals: the ability of anankastic antecedents to supply a goal to (and override competing goals within) the ordering source with respect to which the matrix clause's prioritizing modal is interpreted (see esp. von Fintel and Iatridou 2005). Since many practical conditionals lack goal-expressing antecedents, this issue is orthogonal, as far as I can tell, to the issues raised by representing the meaning of practical conditionals in general.

psychologically, etc.) necessary condition for another state of affairs that is regarded, in a loose sense, as an end, goal, or *telos*.

- (10) a. If the boat is to move, the anchor must be raised
 b. \approx To move the boat, the anchor must be raised
- (11) a. If the clock is to read 12, it must first read 11
 b. \approx For the clock to read 12, it must first read 11

Obviously these sentences have nothing to do with preferability, conditional or otherwise.¹² Anankastic conditionals (and conditionals generally) that express pure relationships of necessity between states of affairs are not correctly understood under the rubric of PCT.

4 Three data points for practical conditionals

If practical conditionals intuitively express the preferability of a state of affairs, conditional on a relevant situation obtaining, it is unsurprising that certain difficulties with the representation of conditional utility re-arise when we shift our focus to them. Conditional preferability is obviously tied rather closely to the notion of conditional utility. Determining the conditional preferability of *B* on *A* is a matter of comparing *B*'s conditional utility on *A* to *C*'s conditional utility on *A* (for all *C* such that *C* is a relevant alternative to *B* if *A*).

So, the question of how to represent the meaning of a practical conditional [*if A*] [*O(B)*] partly parallels the question about how to represent the conditional utility of *B* conditional on wanting *A*. This section gives three preliminary illustrations of this claim (meanwhile also describing three criteria according to which a reading of a conditional of the form [*if A*] [*O(B)*] may be classified as practical).

¹² Nor do they exhibit some of the other defining characteristics of practical conditionals: (i) they arguably can be represented with logical forms in which a necessity modal takes scope over a conditional (i.e., as a kind of strict conditional); (ii) they validate Genuine Modus Ponens (Sect. 4.2). How to draw the line between anankastics that express conditional preferability and anankastics which express pure relationships of necessity is a tricky question, as is the question of how to derive their very different meanings. We will generally consider only anankastics with a clearly practical meaning while avoiding anankastics expressing pure relationships of necessity. We will also sidestep the vexed question of whether a generalization of the semantics we go on to propose could be extended to account for anankastics expressing pure relationships of necessity. (I think it can, but that is another paper entirely.)

On a similar note, even when the modal *O* has a prioritizing interpretation, many conditionals of the form [*if A*] [*O(B)*] do not function to express preferability, conditional or otherwise. Instead they express some sort of law-like relationship between the truth of the antecedent and that of the consequent. For instance:

- (12) a. If stealing is wrong, you shouldn't get your friend to steal
 b. \approx If stealing is wrong, it is wrong to get your friend to steal

It is precisely these sorts of law-like relationships that quantificational approaches to conditionals (the strict conditional and variably strict accounts, most obviously) were designed to handle. I think they do a largely fine job with them. (For a recent argument for the strict conditional view, see, e.g., Gillies 2010.)

Practical conditionals, I will be arguing, do not express a law-like relationship between antecedent and consequent. One virtue of the semantics I defend is that it allows us to give a precise account of the difference between law-like and practical readings of conditionals of the form [*if A*] [*O(B)*]. Stay tuned.

4.1 Scope

Something much like Weirich's argument against reducing conditional to unconditional utility extends to an argument against treating conditionals that intuitively express conditional preferences in terms of preferences that some conditional be made the case. Consider the following case. (I have changed the Cookie case for expository reasons.)

Case 2.

It would, let us suppose, be best (for you) if you wanted a spouse: you should want a spouse. Conditional on wanting a spouse, you should see to it that you have one. But you don't know whether you in fact do want a spouse: your desires concerning this are not transparent to you. Indeed, if you don't actually want a spouse, having one will end up making you fantastically unhappy.

Let Sp be the proposition *you have a spouse*. In Frankfurt's (1971) classic terminology, you have a second-order desire (wanting to want Sp), but are uncertain whether that second-order desire is realized, i.e., whether you have the desired first-order desire (wanting Sp). Because pursuing Sp without knowing whether you want Sp is risky, pursuing Sp is irrational.

In such a case, the following things seem to hold of you. (Here and throughout, *should* expresses preferability, in view of what the specified agent wants.)

- (13) a. You should want $Sp \approx O(\textit{want}(Sp))$
 b. If you want Sp , you should see to it that $Sp \approx ???$
 c. You shouldn't (all things considered) see to it that $Sp \approx O(\neg Sp)$
 (Since Sp has high negative relative goodness, conditional on *not* wanting Sp)

How to represent the logical form of practical conditional (13b)? This is a tricky matter: some natural options fail, for reasons we will soon see.

Intuitively, (13b) expresses that Sp is preferred (i.e., to $\neg Sp$), given that you want Sp . There are broadly two kinds of option here for representing the logical form of (13b)¹³:

¹³ Actually, there are three. As is standard among semanticists (cf. Kratzer 1981, 1991), we could (i) treat O as a dyadic generalized quantifier, with both restriction and scope arguments, (ii) treat indicative antecedents as devices for restricting the domains of such quantifiers. Why don't I consider Kratzer's analysis?

First, to handle anankastics, von Stechow and Iatridou (2005) propose a version of the Kratzer semantics that is a version of a strict conditional semantics for practical conditionals (what they call a "nested modal" analysis). Since the clearest way of making a Kratzer-style account work for practical conditionals (theirs) yields a version of the strict conditional analysis of practical conditionals, giving that account separate attention serves no clear purpose here. The entire point of the paper is to argue that strict conditional accounts of practical conditionals are not viable.

Second, many semanticists see Context-Shift/Dynamic accounts of indicative conditionals as the best implementation of Kratzer's central claim: that the function of an *if*-clause is to restrict. A prime example is the Dynamic account of conditional questions due to Isaacs and Rawlins (2008), on which indicative antecedents function to *restrict the force* of the matrix clause to a derived context, rather than the domain of a quantifier. Context-Shift/Dynamic accounts, which are a major focus of this paper, can, in this sense, be regarded as the best representatives of Kratzer's most important idea.

- **Wide-scoping.** The logical form of (13b) is $O(\text{want}(\text{Sp}) \rightarrow \text{Sp})$.
- **Narrow-scoping.** The logical form of (13b) is $[\text{if } \text{want}(\text{Sp})][O(\text{Sp})]$. ($[\text{if } \phi][\psi]$ is a regimented indicative conditional whose antecedent is ϕ and consequent is ψ .)

Wide-scoping in this case runs into roughly the same sort of problem that afflicted Proposal 1 about conditional utility. (13b) sounds true. But $O(\text{want}(\text{Sp}) \rightarrow \text{Sp})$ must, given our assumptions, be false. Why? Suppose that $O(\text{want}(\text{Sp}) \rightarrow \text{Sp})$ were true. If O expresses a kind of modal necessity, as I will assume here, then the following is a substitution instance of Modal Axiom **K**.

$$O(\text{want}(\text{Sp}) \rightarrow \text{Sp}) \rightarrow (O(\text{want}(\text{Sp})) \rightarrow O(\text{Sp}))$$

Given that $O(\text{want}(\text{Sp}))$ is true, $O(\text{Sp})$ follows, assuming $O(\text{want}(\text{Sp}) \rightarrow \text{Sp})$. This argues against Wide-Scoping the operator O in rendering the logical form of (13b).¹⁴

The parallel with Weirich's trouble case for Proposal 1 is now fully explicit. In both cases, a reductive maneuver is suggested: with conditional utility, that a conditional utility is just the utility of the corresponding conditional, while, with practical conditionals, that a conditional preference for possibilities is equivalent to an unconditional preference for possibilities satisfying the corresponding conditional. But it is not difficult to contrive circumstances in which the *corresponding conditional has negative relative utility* or is *strictly dispreferred*, even when the relevant conditional utility is intuitively positive and the relevant conditional obligation intuitively is the case. So conditional utility does not reduce to the unconditional utility of a conditional, and conditional preferability of some state of affairs cannot be understood as preferability simpliciter of some conditional state of affairs.

Wide-scoping fails, we might say, for the same reason that reductionism for conditional utility functions fails: both fail, in the same way, to appreciate the role that indicative supposition plays in the interpretation of expressions expressing conditional utility or preferability. Indicative supposition can alter one's preferences: to say that A is preferred, given B , should not be taken to tell us very much at all about what is actually preferred.

Footnote 13 continued

To fill this point in a bit more: most of the data that is taken to motivate the Kratzer semantics over competitors is equally well-handled by Context-Shifty or Dynamic accounts. Further, a purely static (i.e., non-Context-Shifty, non-Dynamic) Kratzer-style semantics is beset by some quite serious difficulties. It is hard to see how do devise a principled Kratzer-style analysis of indicatives (i) whose matrix clause is, for instance, a conjunction of two distinct modalities (*John might be at the party, but Sally can't be*) or a conjunction of a modality with a non-modal sentence (*John might be at the party, but Sally isn't*) (Gillies 2010: 4:33), (ii) with clearly non-quantificational matrix clauses, e.g., indicative conditional imperatives (Charlow 2011, 2013b) and indicative conditional questions. (The difficulties a static Kratzer-style semantics encounters with practical conditionals are, in my view, to be regarded as another data point in the case against it and in favor of the Context-Shifty or Dynamic understanding of indicative conditionals.)

¹⁴ Assuming O , whatever its semantics, validates **K** is plausible. But the argument against Wide-Scoping does not require this. As with the case that made trouble for Proposal 1, we might suppose that you take notice of the risk of acquiring a family, and decide (correctly) against it. Given what you know, the only way to make it the case that $(\text{want}(\text{Sp}) \rightarrow \text{Sp})$ is not to want Sp . But that conflicts with our supposition that you should want Sp . So, $O(\text{want}(\text{Sp}) \rightarrow \text{Sp})$ is actually, given what you know, false. But, intuitively, the corresponding practical conditional is true. So the logical form of the latter cannot be given by the former.

More constructively, a necessary condition on whether a reading of a given conditional $[if\ A][O(B)]$ is practical is its resistance to representation using a wide-scoped logical form. This, in turn, can generally be tested by considering whether $[if\ A][O(B)]$ and $O(A)$ together *entail* $O(B)$. If so, the reading of the conditional under consideration is generally not practical, since the conditional generally does not express a suppositional preference. For instance, recall (2), repeated here:

(14) If you want to go to Harlem, you should take the A-train

Suppose this is true—it is solid advice. Suppose also that you *should* want to go to Harlem. Does it *follow*, as a matter of logic, that you should take the A-train? Of course not. If you *don't* want to go to Harlem—if, say, you want to remain in Brooklyn—taking the A-train is potentially strictly dispreferred.¹⁵

4.2 Bootstrapping

Practical conditionals, on their intended readings, do not generally license what I will dub, following, e.g., Broome (1999, 2001), Bootstrapping.

Bootstrapping

Whenever $[if\ want(\phi)][O(\psi)]$ holds, and the agent *actually does want* ϕ , the agent *actually ought* to see to it that ψ .

Recall the Spouse case. Given Sp's high *negative* comparative desirability, conditional on the agent not wanting Sp, it seems the agent—who, while aware of these things is also basically unsure whether she in fact wants Sp—actually ought *not* to see to it that Sp.

I don't wish to deny the existence of a reading of practical conditionals like (13b) on which they license bootstrapping. (Nor do I wish to affirm the existence of such a reading.) I insist only that there is also a (clear, salient) reading on which they do not. That, I wish to suggest, is the actual *conditional preference reading*, in Weirich's sense: the reading according to which Sp is preferred, on the supposition that *want*(Sp).

Here is a related way of making the point. Endorsing bootstrapping means predicting any instance of the following set of sentences *inconsistent*, even when you have no idea whether you actually want ϕ . (Small caps indicate intended stress.)

- (15) a. IF you want ϕ , THEN you should see to it that ψ .
 b. BUT, all things considered, you shouldn't see to it that ψ .
 c. You do want ϕ .

But there is, I insist, a suppositional reading of (15a) on which it is true, even when (15b) is true and you happen to want ϕ . What you should do, given that you want ϕ , concerns what is preferable from the standpoint of the supposition that you do want ϕ .

¹⁵ Broome (1999, 2001) argues that many indicative conditionals expressing what he terms “normative requirements” are ripe for representation using wide-scope logical forms. If he is right, then the relevant readings of these conditionals are ruled out as practical according to this criterion.

What you should actually do concerns what is preferable from the standpoint of your actual information. Again, this is not to deny that there is a reading of (15a) on which it is incompatible with (15b), when it happens that (15c).¹⁶ It is to insist that there is what we might call a “deliberative” *should*: one tracking preferability from the standpoint of whatever information is relevant to a given episode of practical deliberation, on which (15a) and (15b) are compatible.

The same phenomenon is what plausibly licenses the following sequences of practical conditionals and bare pieces of practical advice.

- (16) a. IF it’s sunny, THEN, of course, you should leave the umbrella.
 b. BUT, since we won’t know until we’re outside, you should bring it.
- (17) a. IF the drug prevents infection, THEN you should take it.
 b. BUT, since no clinical studies have shown this, it’s better that you don’t.

If Bootstrapping is valid, then the second utterance in these sequences should not be licensed: the speaker is saying something that may, for all she knows, be false.

Supposition and Bootstrapping. What would be expressed by a conditional preference reading of (13b) is clear enough: from the standpoint of the supposition that you want Sp, Sp is preferred. If there is such a reading available for (13b) (as, I maintain, there is), we would not actually expect that it would license bootstrapping. What an agent should do—at least in the sense relevant to rational decision-making—depends on her information. When something is preferred on *supposition* that ϕ , this does not generally imply anything about what is preferred *simpliciter*, even if ϕ is true. For the actually relevant information—the information available to the agent in her decision situation—may be compatible with both ϕ and $\neg\phi$. A situation in which an agent knows the relevant condition obtains is seldom actual.

In sum, supposing (i) the practical conditional [*if want*(Sp)][*O*(Sp)] expresses conditional preferability, (ii) conditional preferability is to be understood as supposi-

¹⁶ What more could we say about such readings? How might we want to generate them? It is my view that the difference between the two readings—if two readings are well-attested—will lie in how their consequent modals are interpreted. “Deliberative” readings of prioritizing modals have a semantics that will be sensitive in various ways to suppositionally introduced information (in particular, whether that information settles how to achieve the relevant goals), while non-deliberative readings of prioritizing modals will have a semantics that is not sensitive to suppositionally introduced information. (See Charlow (2013a) Cariani et al. (2013), for discussion and a semantics for the relevant modals that makes good on this rough sketch.) The reason I hedge about allowing such readings is that I am inclined to think that indicative conditionals have readings along the lines suggested by the Ramsey Test—an indicative is evaluated by adding the antecedent to one’s suppositions and evaluating the consequent from the standpoint of that supposition. The procedure presupposes, in a loose (but precisifiable) sense, that the information borne by the antecedent is treated as *relevant* to the evaluation of the consequent. It would, in my view, be at least somewhat surprising to observe indicative conditionals whose consequent modals were insensitive to suppositionally introduced information. Maybe the conditionals mentioned in Footnote 12 satisfy this condition. Maybe the readings become easier to access when the relevant modality is strong, rather than weak (cf. Charlow 2013a, pp. 2305–2306). I am unsure. I am quite sure, though, that the readings I appeal to in the main text are real.

tional preferability, we will expect such a conditional to license a merely ersatz kind of bootstrapping.

Ersatz Bootstrapping

Whenever *[if want(ϕ)]**[O(ψ)]* holds, and she *accepts* or *comes to accept want(ϕ)*, she ought to see to it that ψ .

This is another key characteristic of practical conditionals: because they express conditional (i.e., suppositional) preferability, and because what is conditionally preferable—indeed, even the criteria according to which states of affairs are ranked according to preferability—may change depending on conditionally introduced information (see, e.g., Charlow 2013a), they go in only for Ersatz Bootstrapping.

Modus Ponens. To better characterize what's going on here, here is some helpful, and largely standard, semantic machinery. (It will be appealed to over the course of the paper.)

The practical modality O (as well as indicatives whose matrix clause expresses such a modality) is standardly interpreted with respect to a *modal base* and some sort of *selection function*, both furnished by some sort of *state* (to be identified, depending on one's metasemantic inclinations, with a context, state of mind, etc.) (see, e.g., Kratzer 1981). The modal base provides the set of possibilities *relevant* at a state, while the selection function selects the best state-relevant possibilities, yielding the modal O 's domain of quantification.

Definition 1 A state S determines a pair $\langle f_S, \sigma_S \rangle$. $f_S = \lambda i. \{j : j \text{ is a relevant possibility at } \langle S, i \rangle\}$ is a **modal base**, supplying the set of S -relevant possibilities (S -relevant information). And $\sigma_S(i)$ is a **selection function** selecting the best possibilities from a set at i .

The notion ' S -best' is standardly understood relative to a weak preference order on possibilities; the S -best possibilities are simply those that cannot be strictly improved on, relative to the preference order. We will assume, as is fairly standard:¹⁷

Realism

$\sigma_S(i)(f_S(i)) \subseteq f_S(i)$ (the preferred possibilities are always relevant)

Definedness

$\sigma_S(i)(f_S(i))$ is defined and non-empty¹⁸

If we are convinced that indicatives whose consequents contain supposition-sensitive syntactic items do not generally license inference of their consequents when

¹⁷ Notational note: when S is representative of a formal object (a pair of intensions), it is important to relativize the extension of this formal object to a world. When S is representative of something concrete (the state of mind of a world-bound agent), this is not important. To keep our notation unified, parameters are assumed to be world-relative; when world-relativity is not made explicit, such relativity is assumed to be irrelevant. So, e.g., $\sigma_S(f_S) \subseteq f_S$ will sometimes be used to abbreviate a claim like $\sigma_S(i)(f_S(i)) \subseteq f_S(i)$, for arbitrary i . This becomes relevant in Sect. 5.3.

¹⁸ Definedness is the Limit Assumption (Lewis 1973). Relinquishing it affects none of my arguments, and would vastly complicate the semantics. See Kratzer (1981) for discussion.

the relevant supposition is not in force—if, that is to say, we go in only for Ersatz Bootstrapping—we will want such indicatives to validate *only* an “ersatz” form of Modus Ponens.

Genuine Modus Ponens (No!)

The truth of $[if \phi][\psi]$ and ϕ at $\langle S, i \rangle$ implies the truth of ψ at $\langle S, i \rangle$

Ersatz Modus Ponens (Yes!)

The truth (or, possibly, acceptance) of $[if \phi][\psi]$ and ϕ at $\langle S, i \rangle$ implies the truth (acceptance) of ψ at $\langle S|\phi, i \rangle$

Note: $S|\phi$ is S updated with ϕ . For concreteness, suppose for now that updating with ϕ restricts the information to possibilities compatible with ϕ , i.e., that $S|\phi = \langle f', \sigma_S \rangle$, where $f' = \lambda i. f_S(i) \cap \llbracket \phi \rrbracket^S$. (Note: $\llbracket \phi \rrbracket^S$ is the set of possibilities i such that ϕ is true at $\langle S, i \rangle$. When ϕ 's interpretation is state-invariant, we generally omit the state superscript.)

While it might seem unacceptably radical to deny Genuine Modus Ponens for even a *limited* class of natural language indicatives,

- Much recent work on the semantics of indicative conditionals suggests otherwise, specifically (especially) in cases, like ours, where the acceptability of the consequent depends on having access to information¹⁹ bearing on the antecedent (see esp. [Kolodny and MacFarlane 2010](#); [Charlow 2013a](#); [Silk 2013](#); [Cariani et al. 2013](#)).²⁰
- “Genuine” Modus Ponens arguably has no claim to this honorific. Modus Ponens is a rule that legitimates a certain inference pattern under certain conditions; it says nothing about whether those conditions concern the truth, as opposed to acceptance, of the premises. (I am grateful to an anonymous referee for making this point to me. Authors who propose understanding Modus Ponens along the lines of “Ersatz” Modus Ponens include [Gillies \(2009\)](#), [Yalcin \(2012a\)](#).)

¹⁹ Some readers will notice that phrases like “has access to the information that the antecedent is true” are not entirely happy when talking about the consequent’s acceptability—they call to mind naïve formulations of the Ramsey Test, with all their myriad difficulties (recall discussion in Section 2). Since there is, as far as I know, no non-clumsy way to express the right idea in natural English, here and throughout I will use the phrase “ S has access to the information that ϕ ” (and cognate phrases, such as “information ϕ is available to S ”) as abbreviations for a somewhat less committed notion. To say S has access to the information that ϕ in my sense, is to say, roughly, that an abstract representation of S 's information entails ϕ even while not (necessarily) entailing the higher-level proposition *that S 's information entails ϕ* . (There are other ways of making sense of what is typically going on, according to fans of the Ramsey Test, when we suppositionally entertain an indicative antecedent. If you happen to favor one of them, please use it instead.)

²⁰ Though this is a complex issue, it will suffice to note that all the arguments against the various strategies for legitimizing detachment in such cases can be adapted to arguments against the kind of Bootstrapping I am considering here.

Notice that the Wide-Scoper must also deny Modus Ponens for the same class of indicatives. The Wide-Scoper renders $[if \textit{want}(A)][O(A)]$ as $O(\textit{want}(A) \supset A)$; but $O(\textit{want}(A) \supset A)$ and $\textit{want}(A)$ do not validly entail $O(A)$. Indeed, in the ethical literature on normative or rational requirements, this is generally regarded as a *virtue* of the Wide-Scope representation of indicatives describing normative requirements (e.g. [Broome 1999, 2001](#)).

The point bearing emphasis is this. A state S 's having some property P when a supposition is in force (when whether S has P is, in part, a function of available information) ought to imply that S has P when the content of that supposition is accepted at S , *not* when the content of that supposition happens just to be the case.

4.3 Non-monotonicity

An account of indicatives is said to *validate antecedent-strengthening* if it predicts:

Antecedent-Strengthening

The truth of $[if \phi][\psi]$ at $\langle S, i \rangle$ implies the truth of $[if (\phi \wedge \chi)][\psi]$ at $\langle S, i \rangle$

Sobel Sequences. A familiar argument against accounts of indicative conditionals that validate antecedent-strengthening appeals to so-called ‘‘Sobel Sequences,’’ like (18).²¹

- (18) a. If Bob goes to the parade, he'll see Nomar.
- b. #So, if Bob goes to the parade and is stuck behind someone tall, he'll see Nomar.

Any account of indicative conditionals which validates antecedent-strengthening will counterintuitively predict the Sobel Sequence (18) valid.

Any analysis of the indicative conditional as a generalized universal quantifier—by this I will mean any analysis which gives $[if \phi][\psi]$ truth- or acceptance-conditions along the lines of *all ϕ -possibilities meeting some condition C meet some condition D* —will validate antecedent-strengthening (since, in view of the following fact, all ϕ -possibilities meeting C meet D only when all $(\phi \wedge \chi)$ -possibilities meeting C meet D).

$$\{i : i \in \llbracket \phi \wedge \chi \rrbracket \text{ and } C(i)\} \subseteq \{i : i \in \llbracket \phi \rrbracket \text{ and } C(i)\}$$

Things are different with Variably Strict accounts of the indicative conditional, since they do not, strictly speaking, treat the indicative conditional as a univocal generalized quantifier (see Stalnaker 1968, 1975; Lewis 1973). Variably Strict accounts assign an indicative of the form $[if \phi][\psi]$ truth conditions along the following lines: $[if \phi][\psi]$ is true at i just if the most-similar-to- i (relatively to a similarity ordering \leq_i) ϕ -worlds are ψ worlds. This, however, does not entail that the most-similar-to- i $(\phi \wedge \chi)$ -worlds are ψ -worlds, since the most-similar-to- i $(\phi \wedge \chi)$ -worlds are not, in general, a subset of the most-similar-to- i ϕ -worlds.

Possibly : $\{j : j \text{ is } \leq_i\text{-minimal in } \llbracket \phi \wedge \chi \rrbracket\} \not\subseteq \{j : j \text{ is } \leq_i\text{-minimal in } \llbracket \phi \rrbracket\}$

²¹ The argument is usually understood to be a constraint on the semantics of subjunctive or counterfactual conditionals. As Williams (2008) notes, Sobel Sequence-like phenomena arise for indicative conditionals as well. As Moss (2012) notes, the phenomena may have slightly different origins; this will not affect the argument of this paper.

If, for instance, χ 's truth would require a major modification of the circumstances at i (while ϕ 's truth requires only a minor modification), we would expect that some of the most-similar-to- i ($\phi \wedge \chi$)-worlds will not be among the most-similar-to- i ϕ -worlds.²²

This is often cited as motivation for the Variably Strict analysis of indicatives over, e.g., the strict conditional analysis (which is a species of the generalized universal quantifier analysis). The Variably Strict account's explanation of this phenomenon is straightforward: the closest possibilities where Bob goes to the parade and gets stuck behind someone tall are not a subset of the closest possibilities where Bob goes to the parade. Indeed, if (18a) is true, the closest possibilities where Bob goes to the parade plausibly include *no possibilities* where Bob also gets stuck behind someone tall.

Another popular explanation of this phenomenon, from fans of generalized universal quantifier analyses of [*if* ·][·], is that indicative antecedents may *raise*, *make salient*, or *make relevant* possibilities that were not raised/salient/relevant at the original context of utterance (see von Fintel 2001; Gillies 2007; Moss 2012). So the sequence in (18) is fine because the possibilities where Bob goes to the parade and is stuck behind someone tall are not initially relevant; an utterance of (18b) makes them relevant. The initially relevant possibilities where Bob goes to the parade (and which meet some condition C) may all meet condition D , since none of these are possibilities where Bob goes to the parade and is stuck behind someone tall. But the possibilities relevant *after* an utterance of (18b) will include some where Bob goes to the parade and is stuck behind someone tall. And it is easy to see how some (indeed, all) of *these* possibilities might fail to meet D .

This account's plausibility (and the implausibility of Variably Strict analyses) is enhanced by the badness of Reverse Sobel Sequences, like (19) (see esp. von Fintel 2001; Gillies 2007).

- (19) a. If Bob goes to the parade and is stuck behind someone tall, he won't see Nomar.
 b. #But, if Bob goes to the parade, he'll see Nomar.

On this account, (19a) raises the possibility that Bob goes to the parade and is stuck behind someone tall, and asserts that, in any relevant such possibility, Bob doesn't see Nomar. So (19a) makes relevant possibilities where Bob goes to the parade and is stuck behind someone tall. Since (19b) asserts that all relevant possibilities where Bob goes to the parade are possibilities where he sees Nomar, (19b) asserts something that cannot be true, given a prior utterance of (19a). Similar reasoning accounts for the badness of the sequence in (20).

- (20) a. Bob may go to the parade and be stuck behind someone tall.
 b. #But, if Bob goes to the parade, he'll see Nomar.

²² Nevertheless, in many cases—cases where the most-similar-to- i ($\phi \wedge \chi$)-worlds intuitively **are** a subset of the most-similar-to- i ϕ -worlds, in which χ 's truth requires no more of a departure from the actual circumstances than ϕ —we will find that the truth of [*if* ($\phi \wedge \chi$)] $[\psi]$ at i is implied by the truth of [*if* ϕ] $[\psi]$ at i . This will be important later on.

Practical Conditionals. Unsurprisingly, we notice failures of antecedent-strengthening with indicative conditionals whose consequents are practical.

- (21) a. If you want to put sugar in your soup, you should be tested for diabetes!
 b. #So, if you want to put sugar in your soup because you want extra fuel for the marathon, you should be tested for diabetes. (Hare 1971)
- (22) a. If you want to go to Harlem, you should take the A-train.
 b. #So, if you want to go to Harlem and pick up a friend on the way, you should take the A-train.
- (23) a. If you want to eat out, you should drive to the suburbs for Chinese.
 b. #So, if you want to eat out and go clubbing afterward, you should drive to the suburbs for Chinese.

In principle, we may explain these failures by appeal to the same resources used to explain Sobel Sequences like (18).²³ But, although, I will argue, this may work for examples like (21), it leaves much to be desired for anankastic conditionals like (22)–(23).

As noted above, anankastic conditionals are generally paraphrasable as (and intuitively equivalent to) “teleological modal[s] with a purpose *to*-clause,” as in (25) and (26) (Sæbø 2001; von Fintel and Iatridou 2005).

- (25) a. To go to Harlem, you should take the A-train.
 b. #So, to get there and pick up a friend on the way, you should take the A-train.
- (26) a. To eat out, you drive to the suburbs for Chinese.
 b. #So, to eat out and go clubbing, you should drive to the suburbs for Chinese.
- (27) a. ??To put sugar in your soup, you should be tested for diabetes.
 b. ??So, to put sugar in your soup because you want extra fuel for the marathon, you should be tested for diabetes.

²³ A worry: Sobel Sequences, as traditionally described, involve observations about consistency relations between pairs of conditionals, one of whose antecedents is a conjunction one of whose conjuncts is the antecedent of the other. One potential reading of a sequence like (22) has it falling into this syntactic pattern. Another—indeed, the more natural rendering of the sentence—does not. Rather the sequence is represented as follows, with the conjunction taking narrow scope with respect to *want*.

- (24) a. If you want to go to Harlem, you should take the A-train.
 b. #So, if you want [to go to Harlem and pick up a friend on the way], you should take the A-train.

Wanting a conjunction does not generally imply wanting its conjuncts; one can want toast and chicken salad, without wanting the toast. So, the antecedent of the second conditional is not, strictly speaking, a strengthening of the antecedent of the first. I will simply skirt this by making exclusive use of cases where the agent’s wanting of the conjunction does intuitively imply that she wants both conjuncts, i.e., cases in which wanting $(\phi \wedge \psi)$ does intuitively imply wanting ϕ and wanting ψ (and vice versa). I will be willfully sloppy about the difference between wanting a conjunction and wanting both conjuncts in the remainder of the paper.

I do not want to claim that this renders conditionals like those in (21) non-practical.²⁴ But I am happy to allow that it renders them *non-anankastic*.

I do not want to take a stand here on the question of whether Sobel Sequences are to be given a semantic account (as in Gillies 2007) or a pragmatic account (as in Moss 2012). What I want to argue is that the failure of antecedent-strengthening for *anankastic* conditionals:

- Is an evidently semantic phenomenon, deserving a semantic explanation
- Does not admit of the same explanation as Sobel Sequences (regardless of whether that explanation is semantic or pragmatic in character)

These points are made in service of a larger point: Sobel Sequences and the failure of left-downward-monotonicity for anankastics are different phenomena, with different explanations. A semantics for anankastics which accounts for their peculiar failure to license antecedent-strengthening is to be preferred.

The Non-Monotonicity of Goal-Realization. I will assume that anankastic conditionals and their teleological modal paraphrases are semantically equivalent. So we can approach the semantics of anankastics indirectly, by exploring what the teleological modal paraphrases *say*. What they say is that the best way of realizing the goal in the subordinated purpose-clause is to perform the action described in the matrix clause.

Of course, a course of action *A* may be the best way to realize a desire *D*, without it being the case that *A* is the best way to realize a *stronger* desire ($D \wedge D'$). Desire-implementation, as we have known since at least the practical syllogisms of Aristotle, thus has a decidedly *non-monotonic* flavor: what is best for a given end may fail to be best for a set of ends properly including that end. Since teleological modals express best-ness, relative to a given set of ends, we would therefore *expect* that:

Goal Non-Monotonicity

$[to \phi][O(\psi)]$ does not imply $[to \phi \wedge \chi][O(\psi)]$ (*nor vice versa*)

On various pictures about the relationship between sets of goals and selection functions, this can be easily accommodated in the semantics for teleological modals. Here is an example familiar from the semantics literature on practical modals. When *G* is a set of goals, Kratzer (1981) holds that a possibility *i* is at least as *G*-good as *j*—notation $i \leq_G j$ —just if everything in *G* satisfied by *j* is satisfied by *i*.

$$i \leq_G j \text{ iff } \{p \in G : j \in p\} \subseteq \{p \in G : i \in p\}$$

The selection function σ_G selects the possibilities (from a domain *D*) that are minimal relative to this ordering; these are the possibilities that are good enough relative to *G*.

$$\sigma_G(D) = \{i \in D : \neg \exists j \in D : j \leq_G i \text{ and } i \not\leq_G j\}$$

²⁴ To be honest, I don't know whether the (21) conditionals express conditional preferability or not. Mostly I just find them very unusual—probably, I would guess, because the *if you want to...* form is tied so closely to an anankastic meaning.

Here are two things to note about this. First:

- What's good enough for achieving ϕ is **not** always good enough for achieving $\phi \wedge \chi$

Explanation: clearly a possibility i that is good enough relative to G can fail to be good enough relative to a strengthened goal set G^+ such that $G \subset G^+$ (since some other j equally G -good to i may satisfy some p such that $p \notin G$ but $p \in G^+$ that i fails to satisfy, so that i ceases to be minimal relative to the ordering characterized by G^+). Taking the A-train may be a good way to get to Harlem, e.g., but a bad way to get to Harlem and pick up a friend on the way. Second:

- What's **not** good enough for achieving ϕ is sometimes good enough for achieving $\phi \wedge \chi$

Explanation: clearly a possibility i that is **not** good enough relative to G can be good enough relative to a strengthened goal set G^+ such that $G \subset G^+$. Suppose, e.g., there is some $p \notin G$ such that $p \in G^+$ that no possibility $j \in \sigma_G(D)$ satisfies. Suppose i satisfies this p and everything in G besides. Then $i \in \sigma_{G^+}(D)$. For instance, taking a cab may be a bad way to get to Harlem, but a good way to get to Harlem and pick up a friend on the way.

I think it is apparent that these facts about the connection between sets of goals and practical reasoning give a satisfying explanation of our wariness about inferences of the form given in (28) and their converses.

- (28) a. $[to \phi][O(\psi)]$
 b. $[to \phi \wedge \chi][O(\psi)]$

It is important to notice that, on the (extremely plausible) explanation under consideration, the invalidity of this inference schema is *not* explained by the contextual salience of the relevant goals (as a function of the content of the *to*-clause). The *to*-clause does not affect interpretation by making possibilities *relevant*. Rather, it supplies the *argument* for the selection function with respect to which the modal O in the matrix clause is interpreted.²⁵

More precisely, suppose we understand, as is fairly standard, prioritizing necessity modals to be universal quantifiers over the domain defined by the modal base and selection function. In other words, $O(\phi)$ is true at S just if the S -preferred possibilities are ϕ -possibilities.

Definition 2 A modal formula of the form $O(\phi)$ is true at $\langle S, i \rangle$ (notation: $\llbracket O(\phi) \rrbracket^{S,i} = 1$) iff $\sigma_S(i)(f_S(i)) \subseteq \llbracket \phi \rrbracket^S$.

The *to*-clause's function seems to be roughly this.

$$\llbracket [to \phi][O(\psi)] \rrbracket^{S,i} = 1 \text{ iff } \sigma_{\{\llbracket \phi \rrbracket^S, \dots\}}(i)(f_S(i)) \subseteq \llbracket \psi \rrbracket^S$$

²⁵ A semantics in this sort of vein is actually more appropriate for strong practical necessity modals. Weak practical necessity modals have a more complicated semantics, making use of ordered sequences of selection functions (as argued in von Fintel and Iatridou 2008). I do my best to leave room for this sort of treatment in what follows.

Informally, $[to \phi][O(\psi)]$ is true at S just if the selection function characterized by the goal ϕ (and perhaps some body of secondary considerations supplied at S , as indicated by the ‘...’; cf. von Fintel and Iatridou 2008) recognizes only ψ -possibilities as sufficiently good. This sort of semantics yields Goal Non-Monotonicity as a *prediction*: possibilities preferred relative to a goal $\{\phi\}$ may cease to be preferred relative to a stronger goal $\{\phi \wedge \chi\}$. Possibilities best relative to $\{\phi\}$ may fail to be best relative to $\{\phi \wedge \chi\}$; hence, $[to \phi][O(\psi)]$ does not imply $[to \phi \wedge \chi][O(\psi)]$.

Such facts, I submit, explain the acceptability of sequences (25) and (26). Likewise for anankastics, given their equivalence to their *to*-clause paraphrases. Considerations of relevance and similarity, while important for Sobel Sequences, are irrelevant here.

This highlights a final key characteristic of practical conditionals (specifically, in this case, practical conditionals with a specifically anankastic meaning): because they express conditional preferability, and because what is conditionally preferable may change depending on which goals or outcomes are supposed to be conditionally preferred, strengthening the goal expressed by the antecedent of a practical conditional (if its antecedent does express a goal) may nullify the preferability of the consequent’s prejacet. Although this characteristic is specific to anankastic practical conditionals, it is a special case of a more general characteristic of practical conditionals mentioned in Sect. 4.2: conditionally introducing the information expressed by the antecedent of a practical conditional can affect the criteria against which states of affairs are ranked, conditionally, according to preferability.

Anankastic Sequences Are Not Sobel Sequences. This argument is, in my view, both intuitive and quite plausible. Here I want to present some linguistic data that independently bolsters it.²⁶

Were failures of *to*-clause-strengthening for teleological modals explained as a reflex of, e.g., the contextual salience of relevant goals, we would predict that reversing the order of the modals should wreck the sequence’s acceptability (as reversing Sobel Sequences wrecks their acceptability). This prediction is not borne out by the data: sequences (29) and (30) are, perhaps surprisingly, *fine*. Similarly for the anankastics they paraphrase, (31) and (32). Making a strengthened goal salient, as in (33) and (34), does not affect this. (Note: I have added some explanations of the second utterance in the sequences to help readers access the relevant readings.)

- (29) a. To get to Harlem and pick up a friend on the way, you should take a cab.
 b. However, to get to Harlem, you really should take the A-train (...Cabs are expensive, after all, and your friend can fend for himself!)
- (30) a. To eat out and go clubbing afterward, you should stay in Brooklyn.
 b. However, to eat out, you really should drive to the suburbs for Chinese. (...Who cares about clubbing, when there is a chance to eat Chinese in Flushing?)

²⁶ I want to flag that some of the intuitions here are disputed. While I am pretty confident about the judgments, it is hard to *argue* for them. The importance of this depends on how persuasive the reader finds the account I have just given of failures of *to*-clause strengthening. If one is not tempted by the analogy to Sobel Sequences, the data here is *gripping*.

- (31) a. If you want to go to Harlem and pick up a friend on the way, you should take a cab.
 b. However, if you want to go to Harlem, you really should take the A-train. (...Cabs are expensive, after all, and your friend can fend for himself!)
- (32) a. If you want to eat out and go clubbing afterward, you should stay in Brooklyn.
 b. However, if you want to eat out, you really should drive to the suburbs for Chinese. (...Who cares about clubbing, when there is a chance to eat Chinese in Flushing?)
- (33) a. I don't know whether you wish to go to Harlem and pick up a friend on the way.
 b. But, to get to Harlem, you really should take the A-train. (...Cabs are expensive, after all, and your friend could fend for himself!)
- (34) a. I don't know whether you wish to go to Harlem and pick up a friend on the way.
 b. Still, if you want to go to Harlem, you really should take the A-train. (...Cabs are expensive, after all, and your friend could fend for himself!)

What explains, e.g., the acceptability of sequences (29) and (30)? Note that the speaker in each case expressly acknowledges the *possibility* that their addressee's goal is to go to Harlem and pick up a friend along the way. So one cannot respond to the data by saying the subordinate clauses of the (b) utterances elliptically express the goal *going to Harlem without picking up a friend*. It is more natural to say, simply, that the consistency of these sequences—and the oddness of the sequences which they reverse—should be explained by Goal Non-Monotonicity (together with the associated semantics for teleological modal constructions). Any further explanation of the behavior of these sequences is otiose. Since we distinguish sequences like (22) and (25) from Sobel Sequences, we have no expectation at all that reversing these sequences as in (29) should result in unacceptability. That expectation is borne out by the data and extends to the other sequences cited above. The acceptability of (31) and (32) is explained by the fact that (29) and (30) are their intuitive semantic equivalents.²⁷ And, finally, the acceptability of (33) and (34) is explained, in part, by the fact that anankastics and their paraphrases are not sensitive to salient epistemic possibilities in the way that non-anankastic conditionals are typically sensitive.

A more thorough diagnosis of these phenomena is beyond this paper's scope. (A plausible conjecture, drawing on von Fintel and Iatridou (2005, 2008), is that when the goal of picking up a friend is not explicitly designated as a "primary" goal, relevant "secondary" goals (e.g., cost) render sub-optimal possibilities where the primary goal is achieved at high cost.) Even so, empirical considerations do seem to support the intuitive considerations of the prior section. There is independent reason for thinking

²⁷ As suggested by Footnote 11, the explanation is partial: as von Fintel and Iatridou (2005) note, the equivalence itself is something that requires explanation. But that is not really my focus in this paper.

the phenomena we notice in anankastic sequences have a different source than Sobel (and Reverse Sobel) Sequence phenomena.

5 Conditional selection functions

In this section, I lay out some natural options for the analysis of practical conditionals. Conditional preferences figure prominently in this discussion. If we suppose Weirich's view that conditional utility is utility under indicative supposition, we should, given the close relationship between practical conditionals and conditional utilities, keep an eye to how well any given proposal for practical conditionals might be regarded as implementing this idea. Indeed, I will suggest, (i) failure to implement Weirich's idea is at the heart of the failure of many accounts of practical conditionals, (ii) the success of the account of practical conditionals that I defend can be attributed to its successful semantic implementation of Weirich's idea.

Having rejected Wide-Scoping, we have apparently committed ourselves to assigning practical conditionals logical forms in which the prioritizing modal O is interpreted *in situ*. But this is well short of an analysis. A complete semantics would provide at least:

- An interpretation for the conditional construction [*if* ·][·]
- An interpretation for the prioritizing modal O
- An explanation of the data points described in the previous section

The guiding idea here will be that progress on all points is possible, if we take seriously the Practical Conditionals Thesis. Insight into the semantics of, e.g., (13b) is possible if we take the following *conditional preference gloss* as our guide: *given your wanting Sp, Sp-possibilities are preferred*. Another, rough way of putting the point: if unembedded prioritizing modalities express characteristics of selection functions (roughly corresponding to preferences), prioritizing modalities embedded in the consequents of indicative conditionals express characteristics of *conditional selection functions* (roughly corresponding to conditional preferences). Different approaches to the semantics of practical conditionals can be understood as different proposals about the nature and behavior of conditional selection functions.

In this section, I will present (and, in the subsequent section, critique) two distinct approaches to the semantics of (13b), each of which improves on Wide-Scoping:

- The *Variably Strict* approach (due originally to Stalnaker 1968; Lewis 1973)
- The *Context-Shifting* approach (which has various sources, but see Karttunen 1973; Heim 1983; Yalcin 2007, 2012a; Gillies 2010; Kolodny and MacFarlane 2010)

We will develop two versions of the Context-Shifting approach: one on which indicatives express universal quantification over a domain of possibilities, but shift the context of interpretation for proposition-expressing consequents (Context-Shifty *Cognitivism*), and one on which indicatives are treated as non-quantificational devices for shifting the context of interpretation for non-proposition-expressing consequents (Context-Shifty *Expressivism*).

5.1 Wide-scoping, again

My first step is to reconceive the wide-scope analysis of practical conditionals as a proposal about the nature of conditional selection functions. (The only objective here is to introduce the notion of a conditional selection function.)

Suppose the PCT is true. Then, supposing the sentences in (13) are true at S , and helping ourselves to some notation of the form $\sigma_S(\cdot|Y)$ (expressing the conditional-on- Y selection function, i.e., the selection function given Y) we have the following properties of σ_S .

- (35) a. $\sigma_S(i)(f_S(i)) \subseteq \llbracket \text{want}(\text{Sp}) \rrbracket^S$ (*want*(Sp)-possibilities are preferred)
 b. $\sigma_S(i)(f_S(i)|\text{want}(\text{Sp})) \subseteq \llbracket \text{Sp} \rrbracket^S$ (Given wanting Sp, Sp-worlds are preferred)
 c. $\sigma_S(i)(f_S(i)) \subseteq \llbracket \neg\text{Sp} \rrbracket^S$ (\neg Sp-possibilities are preferred)

The conditions on σ_S found in (35) are, I claim, plausible representations of the satisfaction-conditions of the obligation-sentences in (13). However, the as-yet-uninterpreted notion of a **conditional selection function** (hereafter CSF), introduced to represent the meaning of the conditional obligation-sentence (13b), cries out for clarification.

The Wide-Scoper/Reductionist has a proposal at the ready: a preference for ψ possibilities, given ϕ , is equivalent to an unconditional preference for $(\phi \rightarrow \psi)$ -possibilities.

Reduction for CSFs

$$\sigma_S(i)(f_S(i)|\phi) \subseteq \llbracket \psi \rrbracket^S \text{ iff } \sigma_S(i)(f_S(i)) \subseteq \llbracket \phi \rightarrow \psi \rrbracket^S$$

The reductive proposal for the conditional selection function cannot, though, be correct. Given (35b) and Reduction for CSFs, we have:

- (36) $\sigma_S(i)(f_S(i)) \subseteq \llbracket \text{want}(\text{Sp}) \rightarrow \text{Sp} \rrbracket$ (*want*(Sp) \rightarrow Sp-possibilities are preferred)

The difficulty here is that (36) cannot hold if (35a) and (35c) hold. It is a consequence of (35a) and (35c) that $(\text{want}(\text{Sp}) \wedge \neg\text{Sp})$ -possibilities are preferred. Whence it follows that $(\text{want}(\text{Sp}) \rightarrow \text{Sp})$ -possibilities must be strictly *dispreferred*.

It is plausible to think that what explains the failure of wide-scoping here is this. (13b) expresses a conditional preference: specifically a preference for Sp, given *want*(Sp). The wide-scooper represents this preference as equivalent to an unconditional preference for possibilities satisfying the corresponding conditional. But it is easy to contrive circumstances in which this conditional is *strictly dispreferred*. In a sense, wide-scoping reifies a mistaken understanding of the relevant conditional state of mind: as conditional probability resists reduction to unconditional probability (Lewis 1976), conditional preference resists reduction to unconditional preference.

5.2 Variably strict

The Variably Strict analysis of Stalnaker (1968), Lewis (1973) for the indicative conditional can be informally glossed thus: $[if \phi][\psi]$ is true just if the *closest* ϕ -possibilities are ψ -possibilities (recall Sect. 4.3). More carefully, the Variably Strict analysis interprets indicatives relative to a similarity ordering \leq_i that is *centered* on a point of evaluation $i \in W$:

Centering

$$j \leq_i i \Rightarrow i = j$$

Centering codifies the intuition that i uniquely closest to itself. Let \rightsquigarrow represent a variably strict conditional operator at logical form. On the Variably Strict analysis, an indicative $\phi \rightsquigarrow \psi$ is true at i iff the \leq_i -minimal ϕ -worlds are ψ worlds. Formally, letting $min_{\leq_i}(X)$ denote the \leq_i -minimal worlds in X :

Definition 3 $\llbracket \phi \rightsquigarrow \psi \rrbracket^{S,i} = 1$ iff $min_{\leq_i}(f_S(i) \cap \llbracket \phi \rrbracket^S) \subseteq \llbracket \psi \rrbracket^S$

(Note: though we have consistently treated truth as relative to points of evaluation, this is the first point where this relativity is meant to play an interesting semantic role.)

The Variably Strict analysis of (13b) comes in two steps:

- Step 1. A rendering of (13b)’s logical form as $want(\mathbf{Sp}) \rightsquigarrow O(\mathbf{Sp})$
- Step 2. A statement of truth-conditions for the embedded modal sentence $O(\mathbf{Sp})$. (Here, the fan of Variable Strict-ness is free to appropriate the semantics of Definition 2.)

How does this fare? Pretty well. Informally, $want(\mathbf{Sp}) \rightsquigarrow O(\mathbf{Sp})$ is true just if, in the closest possibilities where you want \mathbf{Sp} , you ought to see to it that \mathbf{Sp} . And this holds just if, in all the closest possibilities where you want \mathbf{Sp} , the preferred possibilities there are all \mathbf{Sp} -possibilities. But, unlike the truth-conditions assigned by the Wide-Scoper, these conditions are completely compatible with the *actually preferred* possibilities being $\neg\mathbf{Sp}$ possibilities, even given an actual preference for possibilities where you want \mathbf{Sp} . $want(\mathbf{Sp}) \rightsquigarrow O(\mathbf{Sp})$ says that in (close) possibilities where you want \mathbf{Sp} , \mathbf{Sp} is preferred. A conclusion about what is *actually preferred* follows when you do want \mathbf{Sp} , but not, in general, otherwise. Recall that uncertainty about whether you want \mathbf{Sp} is built into (indeed, integral to the plausibility of) the case. So it is indeterminate, given the description of the case, whether \mathbf{Sp} is actually preferred. And that, according to the fan of Variable Strictness, is as it ought to be. (There is some *serious* fudging going on here, as we will see in Sect. 6.1.)

Unlike the Wide-Scope analysis, then, the Variably Strict analysis invalidates the move from (13a) and (13b) to $O(\mathbf{Sp})$. In light of PCT, we may attribute this to its proprietary understanding of CSFs. On the Variably Strict analysis, what is preferred, given ϕ , is to be identified with what is preferred throughout the relevant ϕ -possibilities.

Preferability-Given- ϕ as Preferability Throughout ϕ -Worlds

$$\sigma_S(i)(f_S(i)|\phi) \subseteq \llbracket \psi \rrbracket \text{ iff } \forall j \in min_{\leq_i}(f_S(i) \cap \llbracket \phi \rrbracket) : \sigma_S(j)(f_S(j)) \subseteq \llbracket \psi \rrbracket$$

The Wide-Scoper’s difficulties traced to her reduction of conditional preference to unconditional preference for a corresponding conditional state of affairs. The Variably

Strict analysis avoids such a reduction: for possibilities to be preferred, given ϕ , is for them to be preferred throughout the modal base, when it has been incremented with the information that ϕ .

The Variably Strict analysis has a prima facie plausible claim to being a semantic implementation of the Weirich-inspired thesis that being preferred, conditional on ϕ , is just being preferred on the supposition that ϕ . Its instruction for evaluating the relevant practical conditional (13b) is something like: consider all the closest possibilities where you want **Sp**. In all such possibilities, is **Sp** preferable? If so, then the conditional is true, otherwise not. On this semantics, the act of supposition is represented as shifting to a domain which entails the antecedent. The act of evaluation, relative to the supposition, is represented as assessing whether **Sp** is preferable in every world in that domain. So far, so good.

5.3 Context-shifting

What I will call the Context-Shifting approach draws on the intuition that checking for the truth of $[if \phi][\psi]$ at a state S involves checking whether ψ holds at the result of updating S with ϕ .²⁸

A variety of approaches in the Context-Shifting vein are possible. All have some prima facie claim to implementing the indicative supposition understanding of conditional preference; whether this claim is substantiated will depend on the details. Different versions of the Context-Shifting approach are obtained by using different understandings of the intuitive notions of:

- What it is for a state to *support* a practical sentence (alternatively, for a practical sentence to *hold at* a state)
- What is the *result of updating* a state with a sentence

Indeed, as we will see, the Context-Shifting approach admits both *truth-conditional* and *non-truth-conditional* understandings of the relevant practical conditionals.²⁹

“Metaphysical” aside. Before describing how Context-Shifting accounts are supposed to work, I want to flag something important. States, as we have defined them, are, formally speaking, *intensional objects*—functions from worlds to extensions. But often we are interested in extensional characteristics of states; for example, regarding

²⁸ For riffs on this theme, see Karttunen (1973), Heim (1983), Yalcin (2007, 2012a), Gillies (2010), Kolodny and MacFarlane (2010). I will be developing the Context-Shifting approach within a dynamic (update-theoretic) framework, but that is not strictly essential to the Context-Shifting story. Indeed, Yalcin (2007, 2012a), Gillies (2010), Kolodny and MacFarlane (2010) show how to develop the same sort of account within a non-dynamic, truth-conditional framework. For most purposes, the choice between dynamic and non-dynamic implementations is irrelevant (and the systems will seem to be quite closely related). I prefer dynamic implementations on the grounds that it is unclear how such an account might be extended to indicative conditionals with non-proposition-expressing consequents—e.g., conditional imperatives.

²⁹ This is fortunate, for it will allow us to generalize our treatment of these conditionals to corresponding conditional imperatives, which exhibit phenomena in all respects similar to the phenomena for the corresponding practical conditionals, but seem not to admit of truth-conditional analysis. It should also be of interest to Expressivists in meta-ethics.

a state as an updatable object (of the sort that is meant to be representative of a concrete agent’s state of mind) means implicitly thinking of it as a fixed extensional object. To finesse this, when treating states as updatable objects (of the sort that are representative of a concrete agent’s state of mind), we will simply assume that we are talking about a concrete state at a given world, and *omit any mention of world-relativity*. Whenever we are treating states as intensional entities, *we will indicate this by making the relevant world-relativity explicit*. As we will go on to see, one version of the Context-Shifting account (Cognitivism) makes use of states in both senses; the other (Expressivism) makes use of states only in the concrete sense.

The formal story. The fundamental semantic notion in update-semantic frameworks is that of a state *supporting* a sentence (alternatively, that of a sentence *holding at* a state). A state S supports a sentence ϕ just if the information borne by ϕ is already borne by S .

Definition 4 A state S supports ϕ (notation: $S \models \phi$) iff $S|\phi| = S$

On the Context-Shifting approach, then, checking for the truth of $[if \phi][\psi]$ involves checking whether ψ holds at a state at which ϕ also holds.³⁰

Context-Shifty Conditionals

$$S \models [if \phi][\psi] \text{ iff } S|\phi| \models \psi$$

Note: it is a consequence of this and Definition 4 that $S \models [if \phi][\psi]$ iff $S|\phi||\psi| = S|\phi|$.

In simple update-semantic systems, a state S is *identified with* its “informational” parameter f_S , and update amounts to monotonic information acquisition, i.e., restriction of f_S . In such systems, the following definition of the dynamic interpretation function $|\cdot|$ suffices.

The Dynamic Interpretation Function $|\cdot|$

$$\begin{aligned} f_S|p| &= f_S \cap \llbracket p \rrbracket \\ f_S|\phi \wedge \psi| &= f_S|\phi||\psi| \\ f_S|\neg\phi| &= f_S - f_S|\phi| \end{aligned}$$

Indeed, in such systems, a recursive definition of $|\cdot|$ is ultimately dispensable, since updating with ϕ amounts just to incrementing the information with ϕ :

$$f_S|\phi| = f_S \cap \llbracket \phi \rrbracket$$

This system must be revised, if we understand the object language as containing practical sentences, i.e., sentences of the form $O(\phi)$. Of course, the definition of $|\cdot|$ must be augmented with a clause for such formulas. Beyond this, we must broaden our understanding of the parameter of the state on which discourse can “bear.” Interpretation of such sentences at a state S is sensitive, not just to the information, but also the selection function σ_S .

³⁰ Strictly, $S|\phi|$ need not support ϕ . Since I deal only with antecedents that are “persistent” (more on persistence below), I will generally ignore this wrinkle. I am inclined, anyway, to insist, following Yalcin (2007), Kolodny and MacFarlane (2010), that the state against which the consequent is checked support the antecedent.

Accommodating practical sentences demands saying how the content of such sentences relates to the selection function. I will consider two conceptions of that relation here. According to Cognitivism, practical sentences bear on one’s *information about* the selection function. According to Expressivism, practical sentences bear “directly” on the selection function—their content is practical, rather than informational, in character. Let’s take these in order.

5.3.1 Cognitivism

According to Cognitivism, the content of a practical sentence is *informational* in nature, in at least two senses. First, a state supports a practical sentence just if the information already entails it. Second, updating with a practical sentence yields a state that supports the sentence.

Practical Sentences Have Informational Content

$$\langle f_S, \sigma_S \rangle \models O(\phi) \text{ iff } f_S \subseteq \llbracket O(\phi) \rrbracket^S$$

$$\langle f_S, \sigma_S \rangle | O(\phi) | = \langle f_S \cap \llbracket O(\phi) \rrbracket^S, \sigma_S \rangle$$

Cognitivists hold that a practical sentence bears information by expressing a proposition. Which proposition? Here the Cognitivist is free to appropriate the standard semantics for O , on which $\llbracket O(\phi) \rrbracket^S = \{i : \sigma_S(i)(f_S(i)) \subseteq \llbracket \phi \rrbracket\}$. (Notice the implicit shift from states as extensional entities to states as intensional entities.) To update on a practical sentence $O(\phi)$, then, is, roughly, to *represent the selection function* as having some property—the property of preferring ϕ .

Combining Cognitivism with the Context-Shifting account of indicatives yields an analysis of the problematic (13b), in two steps:

Context-Shifty Cognitivism (CSC)

- Step 1. Context-Shifting renders the logical form of the conditional *ought* as the regimented indicative conditional $[if\ want(\text{Sp})][O(\text{Sp})]$, yielding the following analysis of (13b):

$$S \models [if\ want(\text{Sp})][O(\text{Sp})] \text{ iff } S|want(\text{Sp})| \models O(\text{Sp})$$

Informally: S supports $[if\ want(\text{Sp})][O(\text{Sp})]$ iff, on updating S with $want(\text{Sp})$, S thereby comes to support $O(\text{Sp})$.

- Step 2. According to Cognitivism, this is the case iff, on updating S with $want(\text{Sp})$, at all possibilities compatible with S ’s (updated) information, Sp -possibilities are preferred.

$$S|want(\text{Sp})| \models O(\text{Sp}) \text{ iff } \forall j \in f_{S|want(\text{Sp})|} : \sigma_{S|want(\text{Sp})|}(j)(f_{S|want(\text{Sp})|}(j)) \subseteq \llbracket \text{Sp} \rrbracket$$

These conditions are quite compatible with the S -preferred possibilities being $\neg\text{Sp}$ possibilities, even given an actual preference for possibilities where Sp is wanted. Suppose, e.g., some possibilities are $\neg want(\text{Sp})$ possibilities: $f_S \cap \llbracket \neg want(\text{Sp}) \rrbracket \neq \emptyset$. In *these* possibilities, recall, $\neg\text{Sp}$ is preferred. So since both $O(\text{Sp})$ and $O(\neg\text{Sp})$ are

true at some relevant possibilities, $f_S \not\subseteq \llbracket O(\text{Sp}) \rrbracket^S$. Hence, according to Cognitivism, $S \not\models O(\text{Sp})$. CSC thus avoids validating the move from (13a) and (13b) to $O(\text{Sp})$, in much the same way as the Variably Strict analysis.

In light of PCT, we can attribute this to its proprietary understanding of CSFs: what is preferred, given ϕ , is identified with what is *known to be preferred* on update with ϕ .

Preferability-Given- ϕ as Known Preferability on Update with ϕ

$$\sigma_S(f_S|\phi) \subseteq \llbracket \psi \rrbracket \text{ iff } \forall j \in f_{S|\phi}(j) : \sigma_{S|\phi}(j)(f_{S|\phi}(j)) \subseteq \llbracket \psi \rrbracket$$

CSC’s conception of conditional preferability is similar to the Variably Strict conception. Both, informally, model conditional preferability as a property of an agent’s *conditional information*: for ψ to be preferred given ϕ is for some ϕ -entailing subset of the agent’s information to entail that ψ is preferred.

5.3.2 *Expressivism*

Another version of the Context-Shifting analysis—owing to a different notion of the relationship between the content of practical sentences and properties of the selection function—is possible. According to Cognitivists, practical sentences have informational content: they express features of one’s conditional preferences by proffering or asserting information about the selection function. But practical sentences might have content that bears on properties of the selection function, without that content being informational in character. Practical sentences might *directly express* properties of a selection function—i.e., manage to express them (in a sense to be precisified) without *asserting* that the selection function has them. Here, briefly, are some examples of what I mean.

Direct Expression: Examples. Similar distinctions have been effectively deployed in the analysis of epistemic modals and ordinary (i.e., non-practical) indicative conditionals. Cognitivist (alternatively known as Descriptivist or Factualist) treatments of epistemic modals assign them propositional content; a sentence like (37) is said to assert that the proposition that it is raining outside is compatible with a relevant body of information.

(37) It might be raining outside.

An alternative idea, which I will call Expressivist,³¹ is that (37) expresses a non-propositional condition such that S supports (37) just if S meets this condition. Specifically, (37) expresses, without asserting, the compatibility of the proposition that it is raining outside with a relevant body of information (Rothschild 2012; Swanson 2006,

³¹ As Yalcin (2011), in particular, has emphasized, the label is a natural one. Expressivism was originally developed as a proposal about the meaning of normative or moral language (Blackburn 1984, 1998; Gibbard 1990, 2003). Such language is claimed to have a peculiar function (being motivational or practical or action-guiding) that cannot be fully accounted for by appeal to its propositional or representational content. Instead, Expressivists claim that such language functions to express properties of motivational or action-guiding states of mind, such that accepting a sentence with normative content is to come to satisfy the property expressed (and thereby to be motivated).

forthcoming; Yalcin 2007, 2011, 2012a). For the Cognitivist, S supports (37) when $f_S \subseteq \llbracket (37) \rrbracket^S$. For the Expressivist, S supports (37) when f_S is compatible with the proposition that it is raining outside: when $f_S \cap \llbracket it \text{ is raining outside} \rrbracket \neq \emptyset$. The latter condition is clearly *non-propositional* in the following sense: satisfying it does not correspond to the state's coming to accept a proposition: there is generally no p such that $f_S \cap \llbracket it \text{ is raining outside} \rrbracket \neq \emptyset$ just if $f_S \subseteq p$.

The dialectic is similar for the analysis of ordinary indicative conditionals. Cognitivists about ordinary indicatives identify a proposition p expressed by a conditional, such that a state S supports the conditional just if $f_S \subseteq p$. Expressivists identify a non-propositional condition expressed by the conditional—e.g., assigning a high conditional probability to the consequent on the antecedent—such that S supports the conditional just if S meets the relevant non-propositional condition (see esp. Swanson forthcoming; Yalcin 2012a). Since there is provably no proposition p such that a high conditional probability is assigned just if a high unconditional probability is assigned to p (Lewis 1976), those who find it fruitful to theorize about the meaning of ordinary indicatives in terms of conditional probabilities will tend toward Expressivism.

This brief introduction to Expressivism hopefully suffices to give a sense of its distinctive semantic methodology.

Expressivism for Practical Modalities. Now, let us return to the question: when does a state support a practical sentence $O(\phi)$? For the Cognitivist, it is for S to satisfy a propositional condition, namely, for the S -relevant information to entail that the selection function has a certain property: the property of requiring ϕ . For the Expressivist, it is for S to satisfy a non-propositional condition. Which such condition? The natural idea is this: S satisfies the relevant non-propositional condition when its selection function simply *has* a certain property, namely, requiring ϕ . (There is no point to relativizing parameters in the Expressivistic definitions to individual points of evaluation. To simplify, we'll decline to do so.)

Practical Sentences Have Preferential Content

$\langle f_S, \sigma_S \rangle \models O(\phi)$ iff $\sigma_S(f_S) \subseteq \llbracket \phi \rrbracket$

$\langle f_S, \sigma_S \rangle | O(\phi) | = \langle f_S, \sigma_{S|O\phi} \rangle$ ($\sigma_{S|O\phi}$ to be defined)

Somewhat impressionistically, the Cognitivist might understand support for practical sentences as analogous to holding a *higher-order* attitude: if σ_S represents the content of an agent's plans (or goals, or desires), the condition of supporting $O(\phi)$ might be glossed as thinking ϕ is a requirement of one's plans (or goals, or desires). The Expressivist, to contrast, might analogize it to a fact about one's *first-order* attitudes: S 's acceptance of $O(\phi)$ is analogous to ϕ just *being* required, in view of the agent's plans (goals, desires). To think $O(\phi)$ is for ϕ to be required, relative to a state representing the agent's information and plans.

Footnote 31 continued

Gibbard, especially, makes systematic use of the distinction between expressing acceptance of a norm and asserting that one accepts a norm. The analysis sketched in this section (and in Sect. 7) can be seen as an attempt to implement the key tenets of Gibbard's expressivism in a formal semantics for natural language. For more on this, see Charlow (2013c).

There are complications in relating this conception of support of a practical sentence to a definition of the update performed by such a sentence on a state—i.e., to a definition of $|O(\phi)|$ —and we will return to this in Sect. 7. But that is not strictly relevant to saying how Expressivist implementation of the Context-Shifting account would go.

To see this consider, once again, the problematic sentence (13b). Combining Cognitivism with the Context-Shifting account of indicatives yields an analysis of (13b), in two steps. Step 1, due to the Context-Shifting part of the account, is unchanged. Step 2, due to the Expressivist part of the account, is novel.

Context-Shifty Expressivism (CSE)

- Step 1. $S \models [if\ want(Sp)][O(Sp)]$ iff $S|want(Sp) \models O(Sp)$
- Step 2. According to Expressivism, this is the case iff, on updating S with $want(Sp)$, Sp -possibilities are preferred.

$$S|want(Sp) \models O(Sp) \text{ iff } \sigma_{S|want(Sp)}(f_{S|want(Sp)}) \subseteq \llbracket Sp \rrbracket$$

According to both Cognitivists and Expressivists, evaluating (13b) is a matter of checking whether the hypothetical state $S|want(A)$ has a certain property. For Cognitivists, that property is *informational*; for Expressivists, it is not.

More generally, given the PCT, Expressivism may be associated with a distinctive understanding of CSFs: what is preferred, given ϕ , can be identified simply—and, it must be said, very intuitively—with what is preferred on update with ϕ .³²

Preferability-Given- B as Preferability on Update with B

$$\sigma_S(f_S|\phi) \subseteq \llbracket \psi \rrbracket \text{ iff } \sigma_{S|\phi}(f_{S|\phi}) \subseteq \llbracket \psi \rrbracket$$

Compare the Cognitivist conception, which understands preferability-given- B as *known* preferability, on update with B . Expressivism, meanwhile, understands conditional selection functions as exactly analogous to conditional utility functions (on Weirich’s suggested understanding): something is conditionally preferred just if it is preferred *under the supposition* that the condition holds—just if it is, in fact, preferred from the vantage of the hypothetical state that results from updating with the relevant condition. Pretty clearly, this proposal also does the job for (13b).

Expressivism, Amended. Expressivism has a minor deficiency which the Variably Strict and Context-Shifty Cognitivist accounts, rather fortuitously, lack. The latter accounts reliably, and correctly, predict the *truth* of $[if\ want(Sp)][O(Sp)]$.

- On the Variably Strict account, $[if\ want(Sp)][O(Sp)]$ is true just if, if i is any sufficiently close possibility where Sp is wanted, the best (in view of what’s wanted) possibilities at i are Sp -possibilities. But, of course, if Sp is wanted at i , then the best (in view of what’s wanted) possibilities at i will typically be Sp -possibilities.

³² In view of Weirich’s Flowers in Bloom example (and complications arising from the Ramsey Test more generally), we will need to introduce a distinction between two ways of updating on a condition—call them “suppositional” and “genuine”—and hold that indicative antecedents induce suppositional, but not genuine, update. For a formal system in which such distinctions are smoothly captured—and which I have elsewhere endorsed (Charlow 2011)—see the Stack-based analysis of Kaufmann (2000). (Thanks to Nathan Howard for discussion here.)

- On CSC, $[if\ want(\mathbf{Sp})][O(\mathbf{Sp})]$ holds at S just if, if i is any S -relevant possibility where \mathbf{Sp} is wanted, then the best (in view of what's wanted) possibilities at i where \mathbf{Sp} is also wanted are \mathbf{Sp} -possibilities. Again, if \mathbf{Sp} is wanted at i , then the best (in view of what's wanted) possibilities at i in which \mathbf{Sp} is wanted will typically be \mathbf{Sp} -possibilities.

Expressivism cannot mimic these predictions directly. Why should updating S on $want(\mathbf{Sp})$ make it the case that the favored possibilities are \mathbf{Sp} -possibilities—why should we expect that $S|want(\mathbf{Sp})|$ favors \mathbf{Sp} ? For the Expressivist, $O(\mathbf{Sp})$ expresses a property of the selection function, rather than information. Supposing that the selection function at S favors $\neg\mathbf{Sp}$, why should the selection function at $S|want(\mathbf{Sp})|$ favor \mathbf{Sp} , given that updating with $want(\mathbf{Sp})$ simply temporarily eliminates $\neg want(\mathbf{Sp})$ -possibilities from relevance?

There are short and long answers to this question. I will sketch the short one here.³³ The short answer is that it is obvious, for independent reasons, that suppositionally updating on $want(\phi)$ will tend to induce a temporary change in the selection function: it makes it the case that ϕ is wanted (relative to the relevant suppositional state).³⁴

We can model this as a kind of general *coherence constraint* on the relationship between one's information about the selection function (understood to include suppositional information) and the properties of the selection function (understood to include its properties under a supposition). Let $\sigma_{\mathcal{R},S}$ be a selection function characterized by a set of propositions $\llbracket\mathcal{R}\rrbracket^S$ —the set of propositions satisfying some relation \mathcal{R} to S (equivalently, a function mapping a proposition to true just if that proposition bears \mathcal{R} to S). So, for instance, $\sigma_{want,S}$ is the selection function characterized by $\llbracket want \rrbracket^S = \{p : p \text{ is wanted at } S\}$. The constraint I have in mind is this:

Coherence

$$\llbracket\phi\rrbracket^S \in \llbracket\mathcal{R}\rrbracket^{S|mathcal{R}(\phi)|}$$

That is to say, if there is a selection function $\sigma_{\mathcal{R},S}$ (whose job is to select the best possibilities, in view of what is \mathcal{R} at S), and if S is updated with the information that ϕ is \mathcal{R} , then the selection function $\sigma_{\mathcal{R},S|mathcal{R}(\phi)|}$ comes to regard ϕ as \mathcal{R} . So updating with $want(\phi)$ makes it the case that the selection function regards ϕ as wanted. So updating S on $want(\phi)$ will typically make it the case that the favored possibilities are ϕ -possibilities.³⁵

³³ The longer answer amounts to an extension of Expressivism to judgments of what is wanted (so that thinking ϕ desired is just to desire ϕ , and determining whether you think ϕ desired is a matter of determining whether you desire ϕ).

³⁴ For a similar take, on which updating on antecedents specifying goals alters the ordering source with respect to which prioritizing modals in the consequent are interpreted, see von Stechow and Iatridou (2005). CSE (plus Coherence) can be easily tweaked in ways that would allow it to replicate von Stechow and Iatridou's analysis of anankastic antecedents as supplying "designated goals." Since the sort of phenomena with which they are concerned are not the focus of this paper, I will not be careful about replicating their analysis here.

³⁵ Coherence constraints must be applied with care, lest we end up making careless predictions about, e.g., the sentences in (21). In this case, the fact that updating on the antecedent *if you want sugar in your soup* makes it the case that the favored possibilities (in view of what you want) are possibilities where you have sugar in your soup does not affect the interpretation of the matrix clause modal. The matrix clause *should* naturally receives an interpretation along the lines of *should, in view of your health*, not *should, in*

Signpost. We sketched three understandings of conditional preferability, each corresponding to a different way of narrow-scoping. Each is associated with an understanding of CSFs that is (i) non-reductive and (ii) plausibly suppositional. Each, therefore, has some prima facie claim to being an implementation of the Weirich-inspired understanding of conditional preferability as preferability under indicative supposition. CSE's claim, perhaps, seems strongest. But so far as the problems confronting Wide-Scoping are concerned, all accounts on a par. In the next section I will try to break the tie for CSE.

6 For CS and E

Variable Strictness, CSC, and CSE all yield prima facie attractive analyses of (13b). In light of the PCT, that is not surprising. All arguably implement, in different ways, Weirich's suggestion that being preferred, conditional on ϕ , is to be preferred on supposition that ϕ . Is there nothing to decide between these approaches?

Here, I will argue that only CSE gets the story right about the data described in Sect. 4. The argument proceeds in two parts. First, I argue that Context-Shifting is the best way around Bootstrapping. Second, I argue that Context-Shifting plus Expressivism is the best way around Non-Monotonicity. Considered together, the cases that distinguish Variably Strict and Context-Shifting approaches (and CSC from CSE) together highlight the superiority of CSE as an implementation of the notion that conditional preferability is suppositional preferability (and ultimately as an implementation of Ramsey's idea that interpretation of $[if \phi][\psi]$ proceeds by evaluating whether ψ from the standpoint of a state that accepts ϕ).

6.1 Bootstrapping

The basic gripe with Variable Strictness is that it turns out to license Bootstrapping:

Footnote 35 continued

view of what you want. Anankastic consequents contain a goal-oriented modal whose goal is *supplied* by their antecedents. The consequents of non-anankastic conditionals with goal-introducing antecedents contain a modal interpreted with respect to a *different goal-set* than that on which the antecedent bears.

What about the meaning of near-anankastics like (38)? (Thanks to a referee for the example.)

(38) If you're averse to being single, you should get a spouse

According to the Coherence constraint, updating on the antecedent of (38) places the proposition that you are single in the set of propositions to whose truth you're averse. What we need to get the right prediction for (38) is for the *negation* of this proposition to be in the set of propositions whose truth you prefer. (I am assuming that the *should* of the matrix clause is expressing necessity with respect to your preferences.) I would suggest that this be handled by adopting a further coherence constraint, such that if p is one of the propositions to whose truth you're averse, then $\neg p$ is one of the propositions whose truth you prefer.

Tweaks on this strategy can be used to handle other near-anankastics like:

(39) If you're afraid of spiders, you should avoid them

(40) If you're attracted to brunettes, you should ask one out

Bootstrapping

Whenever $[\textit{if want}(\phi)][O(\psi)]$ holds, and the agent *actually does want* ϕ , the agent *actually ought* to see to it that ψ .

This is a corollary of the fact that variably strict conditionals validate ‘‘Genuine’’ Modus Ponens.

Claim If $\llbracket \phi \rightsquigarrow \psi \rrbracket^{S,i} = \llbracket \phi \rrbracket^{S,i} = 1$, then $\llbracket \psi \rrbracket^{S,i} = 1$

Proof Suppose $\llbracket \phi \rightsquigarrow \psi \rrbracket^{S,i} = \llbracket \phi \rrbracket^{S,i} = 1$. Then, by Def. 3, $\min_{\leq_i} (f_S(i) \cap \llbracket \phi \rrbracket^S) \subseteq \llbracket \psi \rrbracket^S$. By Centering, $i \in \min_{\leq_i} (f_S(i) \cap \llbracket \phi \rrbracket^S)$. So $i \in \llbracket \psi \rrbracket^S$.

So, on a Variably Strict analysis of (13b), it is immediate that the truth of (13b) and *want*(Sp) implies the truth of $O(\text{Sp})$. So the Variably Strict analysis validates Bootstrapping. But, as we argued in Sect. 4.2, Bootstrapping does not seem to be licensed in case (13). So Variable Strictness cannot be the correct account of this class of practical conditionals.

Fans of Variable Strictness will (and must) react to this observation by insisting on Bootstrapping. There are two ways for them to do this.

- Deny the PCT: while conditional preferability-expressing constructions would not license Bootstrapping, (13b) does not express conditional preferability.
- Embrace the PCT, but deny the suppositional understanding of conditional preferability: while there is a reading of (13b) on which it expresses conditional preferability, conditional preferability is not to be understood as preferability under a supposition.

Both reactions are, I’ll take it, *prima facie* unpalatable. Whether or not the PCT is true, (13b) does not seem to go in for Bootstrapping. (Of course, the PCT explains this, and this is as good a reason as any to take the PCT to be true.) If the PCT is true, then the appeal of Variable Strictness will depend on the extent to which it implements the notion of conditional preferability suggested by the discussion of conditional utility in Sect. 3. Conditional preferability, on this understanding, is preferability under supposition. Preferability under supposition, together with the bare truth of the supposition, simply does not imply preferability simpliciter.

We may discover that appearances are misleading here. For present purposes, however, I take it as given that things are more or less as they appear. We want an analysis of practical conditionals that does justice to (i) the intuitive equivalence of practical conditionals and corresponding statements of conditional preferability and (ii) a suppositional understanding of conditional preferability. Variable Strictness is not up to the task.

Context-Shifting to the Rescue. Context-Shifty accounts do not validate bootstrapping. Supposing $S \models [\textit{if want}(\text{Sp})][O(\text{Sp})]$, whether Sp actually happens to be wanted is, in an important sense, *irrelevant* to whether Sp-possibilities are preferable. There is, of course, a sense in which whether Sp actually happens to be wanted is relevant to whether Sp-possibilities are preferable. To clarify these senses, distinguish *factual* and *informational* relevance.

Factual Relevance

ϕ is factually relevant to the evaluation of $[if \phi][\psi]$ at $\langle S, i \rangle$ just if $\llbracket \phi \rrbracket^{S,i} = 1$ implies that $\llbracket \psi \rrbracket^{S,i} = 1$ (when $\llbracket [if \phi][\psi] \rrbracket^{S,i} = 1$)

Informational Relevance

ϕ is informationally relevant to the evaluation of $[if \phi][\psi]$ at S just if ψ holds at S , given that S supports or accepts ϕ (when S supports or accepts $[if \phi][\psi]$)

Construing the antecedents of practical conditionals as factually relevant means tolerating genuine bootstrapping. Context-Shifty accounts construe such antecedents as informationally relevant; hence they (correctly) tolerate ersatz bootstrapping.

Claim If $S \models [if \phi][O(\psi)]$ and $S \models \phi$, then $S \models O(\psi)$.

Proof Suppose $S \models [if \phi][O(\psi)]$ and $S \models \phi$. Then $S|\phi| \models O(\psi)$. Since $S \models \phi$, $S|\phi| = S$, by Definition (4). Hence $S \models O(\psi)$.

But they fail, in the absence of further assumptions, to tolerate genuine bootstrapping.

Claim $S \models [if \phi][O(\psi)]$ and $\llbracket \phi \rrbracket^{S,i} = 1$ does not imply $S \models O(\psi)$ or $\llbracket O(\psi) \rrbracket^{S,i} = 1$. The claim is obvious and part of the design of Context-Shifting accounts (cf. Gillies 2010).

6.2 Non-monotonicity

While Bootstrapping presents a clear problem only for Variably Strict accounts, accommodating Non-Monotonicity turns out to present a problem for Variable Strictness and CSC (but not CSE).

Against Cognitivism. The conditional of CSC behaves very much like a strict conditional. In particular, it almost immediately validates antecedent-strengthening. This is a corollary of a more general fact: support itself, for CSC, is a monotonic relation (at least when we restrict ϕ and ψ to non-conditional sentences).

Monotonicity of Support

If $S \models \phi$, then $S|\psi| \models \phi$

There are two things to show here. First, that support is monotonic, in the above sense, for CSC. Second, that this leads CSC to validate antecedent-strengthening.

Claim CSC endorses Monotonicity of Support.

Proof Let ϕ, ψ be unconditional sentences, and suppose that $S \models \phi$. Then $S|\phi| = S$. Notice that, for the Cognitivist, $S|\phi| = \langle f_S \cap \llbracket \phi \rrbracket^S, \sigma_S \rangle$. Hence, $f_S = f_S \cap \llbracket \phi \rrbracket^S$. Notice also that $S|\psi| = \langle f_S \cap \llbracket \psi \rrbracket^S, \sigma_S \rangle$. Since $f_S = f_S \cap \llbracket \phi \rrbracket^S$, $f_S \cap \llbracket \psi \rrbracket^S = f_S \cap \llbracket \psi \rrbracket^S \cap \llbracket \phi \rrbracket^S$. Hence $S|\psi| = \langle f_S \cap \llbracket \psi \rrbracket^S \cap \llbracket \phi \rrbracket^S, \sigma_S \rangle$. Hence $S|\psi| = S|\psi||\phi|$. Hence $S|\psi| \models \phi$.

Claim Monotonicity of Support implies Antecedent-Strengthening.

Proof Note that, for any Context-Shifting account, $S \models [if \phi][\psi]$ iff $S|\phi| \models \psi$. Suppose support is monotonic. Then if $S|\phi| \models \psi$, then $S|\phi||\chi| \models \psi$. In which case, since $S|\phi||\chi| = S|\phi \wedge \chi|$, $S|\phi \wedge \chi| \models \psi$. In which case, $S \models [if (\phi \wedge \chi)][\psi]$

In short: CSC problematically validates antecedent-strengthening. Of course, fans of Context-Shifting accounts are aware of this fact, due to (Reverse) Sobel Sequences, and have offered additions to the basic Context-Shifting story to handle it (see esp. Gillies 2007). That is all to the good, so far as (Reverse) Sobel Sequences are concerned. But so far as the sort of non-monotonicity observed in anankastic conditionals is concerned (Sect. 4.3), it is irrelevant.

Why? As argued in Sect. 4.3, the phenomenon that underlies (Reverse) Sobel Sequences is distinct from the phenomenon that underlies the relevant failures of antecedent-strengthening in anankastic conditionals. The latter is due, roughly, to (i) the non-monotonicity of goal-realization, (ii) the ability of the antecedents of anankastic conditionals to *supply goals* relative to which the prioritizing modals in their consequents are interpreted. The former is due to the ability of indicative antecedents to make possibilities *relevant*. The ability of CSC to handle (Reverse) Sobel Sequences indicates nothing about whether they will be able to handle failures of antecedent-strengthening in anankastic conditionals. Indeed, there is at least some reason to think that applying the machinery used to account for (Reverse) Sobel Sequences will make the wrong predictions about anankastic sequences (examples 29–34).

Against Variable Strictness. Perhaps the fan of Variable Strictness is enjoying this turn of events. Whereas she predicts the facts about Sobel Sequences with ease, she also thereby (and problematically) predicts the consistency of Reverse Sobel Sequences like (19). In the case of anankastic sequences, however, there is reason to think that this package of predictions actually squares with the data (see again examples 29–34)!

That is cold comfort. First, one doubts whether the applicability of an explanation originally intended for Sobel Sequences to an unrelated phenomenon should be regarded as good news. Second, and more decisively, the Variably Strict analysis' account of Sobel Sequence (18) originates in a very context-specific fact: the closest possibilities where Bob goes to the parade plausibly include *no possibilities* where he also gets stuck behind someone tall. This follows from the fact that, if it were false, we would expect some of the closest possibilities where Bob goes to the parade (possibilities where he gets stuck behind someone tall) to be possibilities where he fails to see Nomar, thus rendering the conditional (18a) unacceptable.

No such fact is reliably reproduced in the case of our anankastic sequences. The closest possibilities may include possibilities where your advisee wants to get to Harlem (but doesn't care about picking up a friend on the way) *as well as* possibilities where your advisee wants to get to Harlem and pick up a friend on the way. Maybe you don't know (and maybe you don't even care!). Even in such a context, the following conditional is manifestly acceptable.

(41) If you want to go to Harlem, you should take the A-train

Variable Strictness lacks an explanation of the acceptability of (41) in the context in question. The closest possibilities where you want to go to Harlem include, by stipulation, *some possibilities* you want to go to Harlem and pick up a friend on the way. In such possibilities, we expect it to be false that you should take a cab. Variable Strictness counterintuitively predicts the unacceptability of (41) in such a context.³⁶

CSE to the Rescue. Does CSE do better? Yes: whether a conditional allows antecedent-strengthening will generally depend on the acceptance-conditions of its consequent.

According to CSE, indicatives with *information-bearing* consequents—consequents that are accepted just if the information entails them, that express a property of the information-state, rather than the selection function—will license antecedent-strengthening. This is a corollary of the result that CSC validates antecedent-strengthening; CSC differs from CSE only in claiming that practical consequents are information-bearing.

However, indicative conditionals with *practical* consequents—which, according to Expressivists, are accepted iff the selection function *has* a certain property—do not, in general, validate antecedent-strengthening. Per the Coherence constraint, repeated here, antecedents of some such conditionals *introduce goals* and thereby modify the selection function.

Coherence

$$\llbracket \phi \rrbracket^S \in \llbracket \mathcal{R} \rrbracket^{S|\mathcal{R}(\phi)}$$

Thus, an antecedent of the form *want*(ϕ) will tend to make ϕ -possibilities (suppositionally) preferred. And an antecedent of the form *want*(ϕ) \wedge *want*(ψ) will tend to make both ϕ - and ψ -possibilities (suppositionally) preferred. χ -possibilities can be preferred in light of a preference for ϕ , but fail to be preferred in light of a stronger (suppositional) preference for both ϕ and ψ (Sect. 4.3). Courses of action good enough relative to a comparatively *modest* goal-set may cease to be good enough relative to a comparatively *ambitious* goal-set.

Notice, further, that this sort of account straightaway predicts a tight relationship between anankastic conditionals and their teleological modal paraphrases. According to CSE, anankastic antecedents modify the set of goals with respect to which the modal in the consequent is interpreted. This is precisely what the *to*-clauses of their teleological

³⁶ Indeed, we could even imagine that the set of closest possibilities where [you want to get to Harlem and pick up a friend on the way] is a subset of the closest possibilities where [you want to get to Harlem]—that supposing you want to pick up a friend on the way to Harlem requires no more departure from how things actually are than simply supposing you want to go to Harlem.

$$\{j : j \text{ is } \leq_i \text{-minimal in } \llbracket \textit{want}(\textit{Harlem} \wedge \textit{pickup}) \rrbracket\} \subseteq \{j : j \text{ is } \leq_i \text{-minimal in } \llbracket \textit{Harlem} \rrbracket\}$$

Even in such a context, there is nothing remotely wrong with advising the A-train as the best way to get to Harlem, but subsequently advising a cab as the best way to do that and pick up a friend on the way. In such a context, however, Variable Strictness incorrectly predicts that the truth of (42a) implies the truth of (42b).

- (42) a. If you want to go to Harlem, you should take the A-train.
- b. But, if you want to go to Harlem and pick up a friend, you should take a cab.

modal equivalents function to do, on the sort of analysis we suggested in Sect. 4.3:

$$\llbracket [to \phi][O(\psi)] \rrbracket^S = 1 \text{ iff } \sigma_{\llbracket \phi \rrbracket^S, \dots}(f_S) \subseteq \llbracket \psi \rrbracket$$

Why does CSE outperform its competitors here? Plausibly, because it implements the understanding of conditional preferability around which we have structured this paper—conditional preferability as preferability under indicative supposition. Conditional preferability means preferability under indicative supposition, and indicative supposition can, to a degree, alter the desires from which practical deliberation, under a supposition, is undertaken. When we suppose we have a goal ϕ , we deliberate as if we do have that goal (i.e., with a strategic eye to its implementation). When we suppose that we have a goal ϕ as well as a distinct goal ψ , we deliberate as if we have both ϕ and ψ as goals (i.e., with a strategic eye to the implementation of *both*). When we relinquish the supposition that we have ψ as a goal, we again deliberate as if we have A as a goal (with a strategic eye to its implementation). We pivot easily and freely between relatively weak sets of goals and relatively strong sets of goals (and back again).³⁷

It is no mystery why CSE works well for conditionals that express suppositional deliberation about what to do. Its structure tightly tracks the structure of such deliberation. Further, if the anankastic conditionals in question express facts about one's preferences supposing one's possession of certain goals, it is no mystery why they should behave differently than, e.g., non-anankastic indicative conditionals with respect to antecedent-strengthening (as suggested by examples 29–34). Anankastics express the acceptance of a consequent relative to a modified *set of goals*, while non-anankastic indicatives express acceptance of a consequent relative to a modified *information state*. The dynamics of information-evolution and the dynamics of goal-evolution are characterized by altogether different processes and rules.

Coherence's Role. There may be a feeling that we are misallocating credit here, to Expressivism, rather than to Coherence. Suppose the fan of CSC takes on Coherence. Call the resulting view *Modified CSC*. According to Modified CSC, (i) antecedents are context-shifters, (ii) a practical sentence is accepted iff the information entails it, (iii) entertaining an antecedent that expresses a goal can alter the selection function for the interpretation of a practical consequent. Doesn't this help CSC with anankastic sequences?

In a sense, yes.³⁸ But in a more important sense, no. Skirting antecedent-strengthening, we've shown, requires making the support relation non-monotonic for practical sentences. However, if practical sentences genuinely express propositions, as the

³⁷ We'll require a story about why goals suppositionally adopted persist in some contexts, but not in others (examples 29–34). Here a stack-based analysis à la Kaufmann (2000), Isaacs and Rawlins (2008) is quite handy: we can say the stack's top element (representing a suppositional state) is popped by default, i.e., unless there is some reason for it not to be.

³⁸ Proof-sketch. Suppose $f_{S|want(\phi)} \subseteq \llbracket O(\psi) \rrbracket^{S|want(\phi)}$. Although $f_{S|want(\phi)||want(\chi)} \subseteq f_{S|want(\phi)}$, and thus $f_{S|want(\phi)||want(\chi)} \subseteq \llbracket O(\psi) \rrbracket^{S|want(\phi)}$, it does not follow from this that $f_{S|want(\phi)||want(\chi)} \subseteq \llbracket O(\psi) \rrbracket^{S|want(\phi)||want(\chi)}$, on assumption that Coherence holds.

Cognitivist insists, the support relation for such sentences *should* be monotonic. Propositional information, once gained, is *preserved* on subsequent propositional additions. (An *addition* is an update that does not force revision of one's priors; in standard Dynamic Semantic systems, an update u is additive at a state S just if $u(S) \neq \emptyset$. All updates under consideration here are assumed to be additive.) Information that can be lost on subsequent additions is not propositional at all.

An important precedent: in Dynamic Semantics, propositional information is generally taken to be persistent, in the following sense (the classic statement of this position is at [Veltman 1996](#), 225ff; for discussion, see [Yalcin 2012b](#), 273ff). (Note: $S \leq S'$ means that S' is an extension of S , i.e., that S' can be reached from S by a series of non-revision-forcing updates):

Persistence

If $S \leq S'$ and $S \models \phi$, then $S' \models \phi$

Non-monotonicity of support for, e.g., sentences expressing “tests”—updates that query a state for a property, returning the state if it has the property, returning the absurd state if it does not—rather than ordinary intersective updates, is generally taken to imply that such sentences do not express propositions. Indeed, as I'll go on to explain in Sect. 7, there is an attractive story about the semantics of practical sentences on which they do serve to express tests concerning (rather than propositions about) features of the relevant body of preferences. If this is right, there is reason to suspect the Cognitivist is not actually entitled to make use of Coherence. Its adoption will imply that support is non-monotonic for sentences that bear propositional information. Coherence should be married to a non-propositional semantics for practical sentences.

Support for practical sentences is not generally persistent for the Expressivist. Here are several examples. Suppose the S -best possibilities are ϕ -possibilities, i.e., $S \models O(\phi)$, and consider the following scenarios (all of which involve additive changes to S).

1. Let the S -best ψ -possibilities be $\neg\phi$ -possibilities. Then $S|\psi| \not\models O(\phi)$.
2. Let the $S|want(\psi)|$ -best possibilities be $\neg\phi$ -possibilities. Then $S|want(\psi)| \not\models O(\phi)$.
3. Let the $S|!\psi|$ -best possibilities be $\neg\phi$ -possibilities. Then $S|!\psi| \not\models O(\phi)$. (! ψ is an imperative, serving to make its prejacent required. More on imperatives below.)

Intuitively, this is the correct result. It is commonplace and appropriate for additions to one's information or goals to affect one's view about what to do (without thereby requiring that one come to treat an earlier opinion on the matter as *mistaken*). Cognitivism about practical sentences, in view of its prima facie commitment to Persistence for practical sentences, has trouble accounting for this sort of platitude. The Cognitivist must construe such apparently additive changes as revisionary—as involving the rejection of a previously held belief about what one's preferences require.³⁹

³⁹ In principle, troubles like these can be avoided if the Cognitivist insists that the proposition expressed by a practical sentence depends on the state (specifically, preference parameter) against which it is evaluated. This is the strategy that [Rothschild and Yalcin \(2012, Sect. 7.4\)](#) refer to as “Information-sensitivity,” i.e., “the idea of treating the proposition expressed by a sentence as a function of a specific feature of the [state].”

Signpost. I have argued for (i) the superiority of Context-Shifting accounts in handling the phenomenon of Bootstrapping, (ii) the superiority of CSE as an analysis of failures of monotonicity in practical conditionals. In both cases, an account's success is arguably attributable to its implementation of the Weirich-inspired understanding of conditional preferability with which we began this paper. Although Bootstrapping does not provide a reason for favoring one Context-Shifting account over the other, (ii) provides a reason for preferring CSE's handling of both the relevant phenomena (since it provides evidence that CSC is working with the wrong notion of conditional preferability, hence that CSC's explanation of the failure of Bootstrapping is no explanation at all). Having argued for CSE, I now want to spend some time *interpreting* it. I will tackle this in the subsequent section.

7 Clarifying expressivism: imperative versus practical conditionals

In this final section, I will describe some reasons for thinking that CSE yields an appealing treatment of the *cognitive profile* and *force* of practical conditionals. I will do this by examining CSE's treatment of the following questions:

The Cognitive Question

How does an agent *decide whether* $[if \phi][O(\psi)]$?

The Update Question

What happens when an agent *comes to accept* $[if \phi][O(\psi)]$? How do agents *update* on such conditionals?

It would be exceedingly natural, given a truth-conditional semantics for practical language, and a quantificational semantics for conditionals, to answer these questions as follows: (i) deciding whether $[if \phi][O(\psi)]$ is deciding whether $O(\psi)$ holds (i.e., that ψ is preferred) throughout a range of relevant ϕ -possibilities; (ii) deciding that $[if \phi][O(\psi)]$ is a matter of *adjusting your information* so that $O(\psi)$ holds throughout the range of relevant ϕ -possibilities. The discussion of Sect. 6 would seem to suggest that this idea, however natural, is *unworkable*.

Footnote 39 continued

The view of epistemic modals defended in Veltman (1996) meets this criterion, since, although epistemic modals express tests, there is always a set of worlds p such that $f_{S|\diamond\phi} = f_S \cap p$, namely either $W = \top$ (if $f_S \cap \llbracket\phi\rrbracket \neq \emptyset$) or $\emptyset = \perp$ (otherwise).

A state-sensitive "proposition" like this is not really worth the name, however: regardless of the condition of the state, the proposition expressed obviously carries no information about the world. (I do not mean to suggest that Rothschild and Yalcin disagree.) Similarly, if the proposition expressed by $O(\psi)$ is preference-sensitive, it will not carry any information about the world. For note that adopting the Coherence constraint implicitly means treating preferences (in the relevant cases) as world-independent: updating on *want*(ϕ) makes it the case that ϕ is wanted at *every* point of evaluation.

Hence, in the cases relevant here, where what is important is what is preferred on update with sentences of the form *want*(ϕ), $O(\psi)$ will express either \top or \perp . In which case Cognitivism Plus Coherence is effectively a notational variant of the Veltman-style test semantics I go on to state in the subsequent section, and Cognitivism Plus Coherence begins to look an awful lot like Expressivism. (At the very least, the supposed informational content of practical modalities begins to seem elusive.)

Thanks to Seth Yalcin for discussion here.

In the remainder of this paper, I want to explore CSE's answer to the Cognitive and Update Questions. According to CSE:

- a. An agent decides whether $[if \phi][O(\psi)]$ by supposing ϕ and seeing, under this supposition, whether ψ is best

That's to say: entertaining a practical conditional, CSE says, is a matter of *consulting one's preferences* (Expressivism), on update with its antecedent (Context-Shifting). This is distinct from consulting one's beliefs about one's preferences (Cognitivism), on update with the antecedent.⁴⁰

- b. Deciding that $[if \phi][O(\psi)]$ is a matter of ending up in a cognitive state that prefers ψ , on supposition that ϕ

That's to say: deciding that a practical conditional is the case is a matter of ending up in a cognitive state characterized by a specific, suppositional preference. But this, I will argue, should be treated as distinct both from (i) deciding *that* there is such a suppositional preference (hence distinct from a Cognitivist account of endorsement of practical conditionals), and (ii) *coming to have* such a suppositional preference. The discussion here will flesh out the content of these claims and develop a formal framework in which they can be represented.

7.1 The cognitive profile of imperative conditionals

Implementing answers (a) and (b) turns out to present a challenge for Expressivism. For it is a natural idea that conditional *imperatives* like (43), which I will represent schematically as sentences of the form $[if \phi][!\psi]$, are interpreted as proposals to establish suppositional or hypothetical preferences—preferences holding hypothetically, under possible future incrementations of one's information.⁴¹

(43) If you want an apple, take one.

On this sort of picture, deciding to accept a conditional imperative is a matter of adopting a suppositional or hypothetical preference. Accepting an unconditional imperative $!B$ is a matter of adopting an actual preference (i.e., a preference given one's actual information) (see esp. [Portner 2004, 2007](#); [Charlow 2013b](#); [Starr forthcoming](#)). Accepting a conditional imperative is, in other words, a matter of accepting an unconditional imperative in a suppositional cognitive state.

The reason this is problematic is that it would seem to suggest that the cognitive profiles of a practical conditional and a corresponding conditional imperative are *to be identified* (at least partially so). But there is good reason to *reject* any such

⁴⁰ To get the hang of the difference, an agent can similarly attend to the content of her perception (say, of a coffee cup) without attending to her beliefs about the content of her perception. Indeed, the belief that she is perceiving a coffee cup is generally formed *on the basis of* her attention to the content of her perception. Similarly, an agent consulting the content of her preferences is generally the basis for her forming beliefs about the content of her preferences.

⁴¹ This dynamic analysis of conditional imperatives is defended in [Charlow \(2011, Sect. 4.5\)](#) and [Starr \(forthcoming\)](#).

identification. While both practical and imperative conditionals can be “accepted by” or “integrated into” a cognitive state, the phenomenon of accepting a practical conditional seems quite distinct from the phenomenon of accepting an imperative conditional. On the Weirich-inspired picture that I favor, accepting a practical conditional $[if \phi][O(\psi)]$ is a natural upshot of:

- Entertaining whether $[if \phi][O(\psi)] \approx$ supposing ϕ and entertaining whether $O(\psi)$
- Deciding that $[if \phi][O(\psi)] \approx$ deciding that $O(\psi)$, on the supposition that ϕ

But accepting a corresponding imperative conditional cannot result from this sort of process. One cannot entertain whether $[if \phi][!\psi]$, since one can simply never entertain whether $!\psi$, regardless of one’s cognitive state. Nor can one decide that $[if \phi][!\psi]$, since one can never decide that $!\psi$.⁴²

(44) *Sally entertained whether clean your room

(45) *Sally decided that clean your room

According to the Expressivist account of the cognitive profile of practical conditionals that I have been entertaining, deciding that $[if \phi][O(\psi)]$ means ending up in a cognitive state that prefers ψ , on supposition that ϕ . But there is good reason to think updating on the corresponding imperative conditional is also a matter of ending up in a cognitive state that prefers ψ , on supposition that ϕ . So, it doesn’t seem that this can be the right story about accepting a practical conditional. Acceptance of a practical conditional is generally *issue-resolving*, in a way that acceptance of an imperative conditional is not—indeed, cannot be. Expressivism as a theory of practical cognitive force seems not to admit a sufficient distinction between practical and imperative cognitive force. Perhaps we should give CSC another look?

This is fair enough, at least as a description of the account as developed to this point. In the next section, we will show how to correct this fault.

7.2 Updates and queries

It is not difficult, if we are willing to be impressionistic, to identify a difference between acceptance of a practical conditional and acceptance of an imperative conditional. That difference, I will suggest, is just this: accepting a practical conditional is a matter of *verifying* (as a result of successfully *querying* or *testing*) current properties of one’s preferences, while accepting an imperative conditional is a matter of *altering* one’s preferences.

Nevertheless, I want to suggest that those states in which $[if \phi][O(\psi)]$ is accepted—those that pass the query or test—are just those in which $[if \phi][!\psi]$ is accepted:

$$S \models [if \phi][O(\psi)] \text{ iff } S \models [if \phi][!\psi]$$

The practical conditional checks the state for just the property that would lead update with the imperative conditional to idle. Executing a query on a state’s hypothetical

⁴² Generally, entertaining whether ϕ and deciding that ϕ are possible only when ϕ is declarative.

preferences, however, is a fundamentally different operation than adjusting the state's hypothetical preferences. Executing a query on a state's preferences can *resolve issues* concerning its properties (since it is inherently sensitive to features of the state), while altering the state's preferences—acting by way of changing the state's preferences—cannot (since it is insensitive to its features⁴³).

Here is a critical distinction. Querying a state for a property P , in the intended sense, is different from *asking whether* the state has P . Likewise, successfully querying the state for P —verifying that it has P —is distinct from deciding *that* it has P . The latter induces a change to one's state (by updating the state's information about itself), while the former simply leaves the state unchanged. The issue *whether* $[if \phi][O(\psi)]$ is resolved by querying a state, without it being the case that properties of the state are what it is at issue. One decides whether $[if \phi][O(\psi)]$ by testing one's hypothetical preferences, but the question of whether $[if \phi][O(\psi)]$ is not a question *about* one's hypothetical preferences. It is, rather, a non-psychological question—one about what is best, on the supposition that ϕ . It is a question that is decided *from* the standpoint of hypothetical preference, rather than a question *about* that standpoint.

Test semantics for epistemic modality. An account in this vein has been developed for epistemic modals by semanticists in the Dynamic Semantic tradition (see esp. Veltman 1996; von Fintel and Gillies 2007). Its motivating intuitions are well-stated by Frank Veltman.

[A]ll you can do when told that it might be the case that ϕ is to agree or to disagree. If ϕ is acceptable in your information state S , you must accept $\diamond\phi$. And if ϕ is not acceptable in S , neither is $\diamond\phi$. Clearly, then, sentences of the form $\diamond\phi$ provide an invitation to perform a test on S rather than to incorporate some new information in it. (1996, p. 229)

I will briefly explain the account here and extend it to practical conditionals in the next section. On the account, the cognitive force of accepting an epistemically modalized clause $\diamond\phi$ is to verify ϕ 's compatibility with the relevant information. Although the semantics does not strictly represent this, it is natural to divide the process of accepting $\diamond\phi$ into two stages:

- Entertaining whether $\diamond\phi$ (=executing the test).
- Deciding that $\diamond\phi$ (=responding to the successful performance of this test, by remaining in a state compatible with ϕ)

On the other hand, to accept $\neg\diamond\phi$ (or to reject $\diamond\phi$) is to respond to the unsuccessful performance of this test, by remaining in a state which fails to accept $\diamond\phi$ (a state incompatible with ϕ).

⁴³ Changes to states often incorporate preconditions or presuppositions, so that update crashes when those preconditions or presuppositions fail to be met by the state. So there is a sense in which preference-changes might be state-sensitive. But these preconditions are best understood as queries of a state: the state-sensitivity of preference-change is due to their incorporation of such queries. Crucially, however, the update that is *proffered* by an imperative—at least on the analyses of Portner, Starr, and Charlow—is a non-querying update.

Formally, tests are traditionally specified as vacuous (or else catastrophic) updates on states: integrating $\diamond\phi$ involves checking to see whether the input-state accepts $\diamond\phi$, returning that state in the event that it does, returning an absurd state otherwise.

$$S|\diamond\phi| = \begin{cases} S, & \text{if } f_S \cap \llbracket\phi\rrbracket \neq \emptyset \\ \emptyset, & \text{otherwise} \end{cases}$$

This is *not* to be interpreted as suggesting that $\diamond\phi$ proffers an update instruction. Given a state as input, $|\diamond\phi|$ returns information about that state—namely, about its compatibility with ϕ —as output. Executing the test yields this information without attempting to induce the state to either (i) *represent* that information as true (else updating with $\diamond\phi$ would yield a change to S , and it does not), or (ii) induce the state to become compatible with ϕ . If S passes the test, S accepts $\diamond\phi$, and remains unchanged. If S fails to pass the test, S rejects $\diamond\phi$ (since accepting it would lead to an absurd state) and, likewise, remains unchanged.⁴⁴

Test semantics for practical sentences. This is the template I want to pursue for practical sentences (and eventually conditionals). A practical sentence $O(\phi)$, like $\diamond\phi$, will test a state for a property: the property of preferring ϕ (or regarding ϕ as possible, as the case may be).

$$S|O(\phi)| = \begin{cases} S, & \text{if } \sigma_S(f_S) \subseteq \llbracket\phi\rrbracket \\ \emptyset, & \text{otherwise} \end{cases}$$

The force of accepting $O(\phi)$, then, is *not* to enforce coordination with respect to the property of accepting $O(\phi)$; this, rather, is the force of accepting the corresponding imperative sentence $!\phi$. It is, rather, to verify that ϕ is preferred. This process, like the analogous process for epistemic modals, is naturally understood in two stages:

- Entertaining whether $O(\phi)$ (=executing the test)
- Deciding that $O(\phi)$ (=responding to the successful performance of this test, by remaining in a state that prefers ϕ)

On the other hand, to decide that $\neg O\phi$ is to respond to the unsuccessful performance of this test, by remaining in a state which fails to prefer ϕ . I emphasize that entertaining whether $O(\phi)$ is distinct from entertaining whether one's cognitive state prefers ϕ . Entertaining whether $O(\phi)$, cognitively, is a matter of executing a test on one's preferences. An agent assesses whether ϕ is best by *consulting* her preferences. Such an assessment does yield information about the agent's preferences, but the agent's inquiry is not necessarily about *this* information. It is typically, rather, about the practical "worldview" that those preferences encode or represent.

⁴⁴ This is contra Yalcin's (2007) suggestion that the force of $\diamond\phi$ is to *induce coordination* on the property of treating ϕ as compatible with the relevant information. Although a state's interpretation of $\diamond\phi$ may *set the stage* for attempts to coordinate on the property of being informationally compatible with ϕ (if, for instance, interpretation brings to light a failure in coordination), there is no proposal that aims, as such, at such coordination *and* which is also part of the conventional force of $\diamond\phi$. The conflation of testing a state S for a property P with proposing that S come to satisfy P is, I will suggest, what is responsible for the running together of imperative and practical force discussed in Sect. 7.1.

There is much more to say about this proposal, but, however it gets said, notice that, contra the suggestion in Sect. 7.1, embracing Expressivism as a theory of practical cognitive force still allows a wide separation between practical and imperative cognitive force. Though deciding that $O(\phi)$, like accepting the corresponding imperative $!\phi$, is a matter of ending up in a cognitive state that prefers ϕ , the former is generally issue-resolving, since it occurs in the context of entertaining whether $O(\phi)$. Accepting a corresponding imperative is not generally issue-resolving, although both updates terminate in a state that prefers ϕ .

Practical Conditionals, Redux. Extending the story to practical conditionals proceeds much as you'd expect. A practical conditional tests a state's hypothetical or suppositional preferences.

$$S[[if\ \phi][O(\psi)]] = \begin{cases} S, & \text{if } \sigma_{S|\phi}(f_{S|\phi}) \subseteq \llbracket \psi \rrbracket \\ \emptyset, & \text{otherwise} \end{cases}$$

Entertaining a practical conditional is a matter of executing a test on one's suppositional preferences. An agent assesses whether a practical conditional holds by *consulting* her suppositional preferences. Such an assessment yields information about the agent's suppositional preferences, but, again, the agent's inquiry is not necessarily about *this* information. As Blackburn and Gibbard have suggested, the agent experiences her inquiry as an inquiry into what is *good or worth pursuing*.

This is a fully compositional story about the cognitive force of practical conditionals. It is a direct consequence of blending a Context-Shift account of indicatives with a Test-Expressing Expressivism for their practical consequents. On this account, though deciding that $[if\ \phi][O(\psi)]$, like accepting the corresponding imperative $[if\ \phi][!\psi]$, is a matter of ending up in a state that prefers ψ , on the supposition that ϕ , deciding that $[if\ \phi][O(\psi)]$ is generally issue-resolving, since it occurs in the context of entertaining whether $[if\ \phi][O(\psi)]$.

8 Conclusion

This paper had many twists and turns. I will try, briefly, to give it a sense of unity. We began with Weirich's reflections on conditional utility. Those reflections supported a specific understanding of conditional utility—as utility under indicative supposition—which in turn suggested a specific understanding of conditional preferability—as preferability under indicative supposition. We tried out four accounts of the semantics of practical conditionals—conditionals that function to express judgments of conditional preferability—and found that the best fit for the relevant notion of conditional preferability was to be found in CSE.

This was a striking result. The only account of indicative practical conditionals that did justice to the data (as well as our decision-theoretic reflections on conditional preference) was *non-quantificational*. The implications for future theorizing about a unified semantics for indicative conditionals, as a class, are difficult to miss.

Acknowledgments Thanks to Simon Charlow, Nathan Howard, and Seth Yalcin. Special thanks to two anonymous referees and to my *L&P* editor Lisa Matthewson for their comments.

References

- Adams, E. W. (1975). *The logic of conditionals: An application of probability to deductive logic*. Dordrecht: D. Reidel.
- Blackburn, S. (1984). *Spreading the word*. Oxford: Clarendon Press.
- Blackburn, S. (1998). *Ruling passions*. Oxford: Oxford University Press.
- Bradley, R. (1999). Conditional desirability. *Theory and Decision*, 47, 23–55.
- Broome, J. (1999). Normative requirements. *Ratio*, 12, 398–419.
- Broome, J. (2001). Normative practical reasoning. *Aristotelian Society Supplementary Volume*, 75, 175–193.
- Cariani, F., Kaufmann, M., & Kaufmann, S. (2013). Deliberative modality under epistemic uncertainty. *Linguistics and Philosophy*, 36, 225–259. doi:10.1007/s10988-013-9134-4.
- Charlow, N. (2011). *Practical language: Its meaning and use*. Ph.D. dissertation, University of Michigan. <http://www.natecharlow.com/work/dissertation.pdf>.
- Charlow, N. (2013a). What we know and what to do. *Synthese*, 190, 2291–2323.
- Charlow, N. (2013b). Logic and semantics for imperatives. *Journal of Philosophical Logic*. doi:10.1007/s10992-013-9284-4.
- Charlow, N. (2013c). The problem with the Frege-Geach problem. *Philosophical Studies*. doi:10.1007/s11098-013-0119-5.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, 68, 5–20.
- Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge: Harvard University Press.
- Gibbard, A. (2003). *Thinking how to live*. Cambridge: Harvard University Press.
- Gillies, A. S. (2007). Counterfactual scorekeeping. *Linguistics and Philosophy*, 30, 329–360.
- Gillies, A. S. (2009). On truth-conditions for *If* (but not quite only *If*). *The Philosophical Review*, 118, 325–349.
- Gillies, A. S. (2010). Iffiness. *Semantics and Pragmatics*, 3, 1–42.
- Hájek, A. (2012). The fall of ‘Adams’ thesis’?. *Journal of Logic, Language, and Information*, 21, 145–161.
- Hare, R. M. (1971). Wanting: Some pitfalls. In *Practical inferences*. New York: Macmillan.
- Heim, I. (1983). On the projection problem for presuppositions. In M. Barlow, D. Flickinger, & M. Wescot (Eds.), *WCCFL 2: Second annual west coast conference on formal linguistics*. Stanford: CSLI Publications.
- Isaacs, J., & Rawlins, K. (2008). Conditional questions. *Journal of Semantics*, 25, 269–319.
- Jeffrey, R. (1983). *The logic of decision* (2nd ed.). Chicago: University of Chicago Press.
- Joyce, J. (2000). Why we still need the logic of decision. *Philosophy of Science*, 67, S1–S13. <http://www.jstor.org/pss/188653>.
- Karttunen, L. (1973). Presuppositions of compound sentences. *Linguistic Inquiry*, 4, 169–193. <http://www.jstor.org/stable/4177763>.
- Kaufmann, S. (2000). Dynamic context management. In M. Faller, S. Kaufmann, & M. Pauly (Eds.), *Formalizing the dynamics of information*. Stanford: CSLI Publications. <http://faculty.wcas.northwestern.edu/kaufmann/Papers/dcm.pdf>.
- Kaufmann, S., & Schwager, M. (2009). A unified analysis of conditional imperatives. In E. Cormany, S. Ito, & D. Lutz (Eds.), *Proceedings of SALT 19* (pp. 239–256). Ithaca: CLC Publications. semanticsarchive.net/Archive/TYZZYyM/.
- Kolodny, N., & MacFarlane, J. (2010). Ifs and oughts. *Journal of Philosophy*, 107, 115–143. <http://johnmacfarlane.net/ifs-and-oughts.pdf>.
- Kratzer, A. (1981). The notional category of modality. In H. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts* (pp. 38–74). Berlin: De Gruyter.
- Kratzer, A. (1991). Conditionals. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: An international handbook of contemporary research* (pp. 651–656). Berlin: De Gruyter.
- Lewis, D. (1973). *Counterfactuals*. Malden: Basil Blackwell.
- Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review*, 85, 297–315. <http://www.jstor.org/pss/2184045>.
- Moss, S. (2012). On the pragmatics of counterfactuals. *Noûs*, 46, 561–586.

- Portner, P. (2004). The semantics of imperatives within a theory of clause types. In K. Watanabe & R. Young (Eds.), *Proceedings of SALT 14*. Ithaca: CLC Publications. <http://semanticsarchive.net/Archive/mJZGQ4N/>.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15, 351–383.
- Rothschild, D. (2012). Expressing credences. *Proceedings of the Aristotelian Society*, CXII, Part 1, 99–114.
- Rothschild, D., & Yalcin, S. (2012). On the dynamics of conversation. Ms. <http://yalcin.cc/resources/DynamicsConv.pdf>.
- Sæbø, K. J. (2001). Necessary conditions in a natural language. In C. Féry & W. Sternefeld (Eds.), *Audiat Vox Sapientiae: A Festschrift for Arnim von Stechow*. Berlin: Akademie Verlag.
- Schwager, M. (2006). Conditionalized imperatives. In M. Gibson & J. Howell (Eds.), *Proceedings of SALT 16*. Ithaca: CLC Publications. <http://user.uni-frankfurt.de/~scheiner/papers/schwagerFEB07.pdf>.
- Silk, A. (2013). Evidence sensitivity in weak necessity deontic modals. *Journal of Philosophical Logic*. doi:10.1007/s10992-013-9286-2.
- Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory* (pp. 98–112). American philosophical quarterly monograph series 2. Oxford: Blackwell.
- Stalnaker, R. (1975). Indicative conditionals. *Philosophia*, 5, 269–286.
- Starr, W. B. (forthcoming). A preference semantics for imperatives. *Semantics and Pragmatics*.
- Swanson, E. (2006). *Interactions with context*. Ph.D. dissertation, MIT.
- Swanson, E. (forthcoming). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25, 221–261.
- von Fintel, K. (2001). Counterfactuals in a dynamic context. In M. Kenstowicz (Ed.) *Ken Hale: A life in language*. Cambridge, MA: MIT Press. <http://mit.edu/fintel/fintel-2001-counterfactuals.pdf>.
- von Fintel, K., & Gillies, A. S. (2007). An opinionated guide to epistemic modality. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 2). Oxford: Oxford University Press. <http://mit.edu/fintel/fintel-gillies-2007-ose2.pdf>.
- von Fintel, K., & Iatridou, S. (2005). What to do if you want to go to Harlem: Anankastic conditionals and related matters. Ms., MIT. <http://mit.edu/fintel/www/harlem-rutgers.pdf>.
- von Fintel, K., & Iatridou, S. (2008). How to say *ought* in foreign: The composition of weak necessity modals. In J. Guéron & J. Lecarme (Eds.), *Time and modality*. Dordrecht: Springer.
- von Wright, G. H. (1963). *Norm and action: A logical inquiry*. London: Routledge.
- Williams, J. R. G. (2008). Conversation and conditionals. *Philosophical Studies*, 138, 211–223. doi:10.1007/s11098-006-9032-5.
- Weirich, P. (1980). Conditional utility and its place in decision theory. *The Journal of Philosophy*, 77, 702–715.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116, 983–1026.
- Yalcin, S. (2011). Nonfactualism about epistemic modality. In B. Weatherson & A. Egan (Eds.), *Epistemic modality*. Oxford: Oxford University Press.
- Yalcin, S. (2012a). A counterexample to modus tollens. *Journal of Philosophical Logic*, 41, 1001–1024.
- Yalcin, S. (2012b). Dynamic semantics. In G. Russell & D. Fara (Eds.), *Routledge companion to philosophy of language* (pp. 253–279). New York: Routledge.