A Buyer's Guide to Conditionals

Daniel Bonevac (bonevac@mail.utexas.edu)
University of Texas at Austin

Josh Dever (dever@mail.utexas.edu)
University of Texas at Austin

David Sosa (david_sosa@mail.utexas.edu)*
University of Texas at Austin

* The authors' names are listed in alphabetical order.

What kind of generalized quantifier is a conditional? Before passing on to the usual display of paradigmatic (non-)inferences, let us reflect more deeply. Our intuitions come in various kinds, and it is important to consider the more volatile ones first, concerning the kind of notion that we are after, before these are drowned in a list of very specific desiderata. Only in the light of such background intuitions, one can take a proper look at more concrete claims of validity or invalidity of conditional inferences. (van Benthem (1984), 309)

1. Reflecting Deeply

So you want to purchase a conditional. It's a buyer's market these days, with more conditionals out there to choose from than you can shake a stick at. In addition to the classic material conditional, you have available the C.I. Lewis strict conditional (in the S1, S2, and S3 models[1]); 26 flavours of the David Lewis counterfactual (Lewis (1973) (including the Stalnaker conditional of Stalnaker (1968) as the special case of system **VCS**); the test conditional of Veltman (1996)'s Update Semantics; the $\mathbf{S4}_{\rightarrow}$, $\mathbf{H}_{\rightarrow}$, $\mathbf{R}_{\rightarrow}$, $\mathbf{E}_{\rightarrow}$, and $\mathbf{T}_{\rightarrow}$ conditionals of relevance logic (Anderson and Belnap (1975)); the dynamic context-shifting conditionals of von Fintel (2001) and Gillies (2007); the premise semantics for counterfactuals via 'lumps of thought' in Krazter (1989); the normality conditional of Asher (1995) and Morreau (1997); the probabilistic conditionals of Adams (1975) and Appiah (1984); the local and global variants of probabilistic conditionals in Kaufmann (2004); the "ifs as restrictors" view of Lewis (1975); indicative and subjunctive conditionals (a distinction first stressed in Adams (1970); biscuit conditionals (first discussed in Austin (1956) and carefully explored in (e.g.) Siegel (2006)); Weatherson (2008)'s "indexical relativism" conditionals; "conditional assertion" conditionals such as those recently defended by Edgington (1995); backtracking counterfactuals (see Lewis (1979)); the branched-time temporal conditionals of Thomason and Gupta (1980); anankastic conditionals, as examined in von Fintel and Iatridou (2004); projective and hypothetical conditionals

[1] Note that the luxury S4 and S5 models are differentiated from the base models only by their inclusion of further constraints on the modal portion of the logic (with S4 adding $\Box p \supset \Box\Box p$ and S5 adding $\Diamond p \supset \Box\Diamond p$). Thus the C.I. Lewis options in the pure conditional logic are restricted to S1-S3. Interesting, none of these pure conditional logics is strong enough to receive a normal modal semantics (each thus requiring the non-normal worlds techniques of Kripke (1965)). The connection between the strict conditional and a conditional expressing covariation in truth value across all possible worlds is thus perhaps less tight than philosophical received wisdom takes it to be. It is observations of this sort that lead us to suspect that conditionals are rather less well understood, even in the more familiar neighborhoods, than is commonly supposed.

(see Dudman (1991)); the relocated conditionals of Bennett (1988); monkey's uncle conditionals (see Sadock (1977); the consequentialist and non-consequentialist conditionals of Dancygier and Mioduszewska (1984); and the helping, hindering, and lemon balm conditionals of Humberstone (2008).

At this point a standard way to proceed, *qua* consumer, with the choice of conditional is via *feature comparison* – one makes a list of important features of each conditional under consideration (paradigmatically, inferential features; but other issues such as fidelity to natural language data on assertability and truth conditions, or integrability into various philosophical projects, may also loom large here), and then selects the conditional which best matches some pre-theoretic view on the ideal feature set. We feel, however, that this methodology is unproductive. Our intuitive grip on the ideal inferential (etc.) features of the conditional is simply too weak for a feature check list of this sort to move the discussion forward. Fans of the material conditional (see, e.g., Jackson (1987)) have shown that, at a minimum, there is considerable difficulty in separating pragmatic from semantic features when attempting to assess the desirability (or undesirability) of inferential features (such as those of the "paradoxes of the material conditional"). Discussions of probabilistic conditionals and of "conditional assertion" views of the conditional have shown that the very question of whether there is *anything* that properly serves as truth conditions for conditionals is a vexed one. And as logical space is increasingly densely colonized by competing views, the requisite sensitivity to the acceptability of (often recherche) inferential forms increases beyond reasonable limits.

Given all of these factors fogging implementation of the standard methodology, our view is that it is premature to have these disputes. Before we can select a conditional based on its inferential (etc.) features, we need to understand further the sources and justifications of these sundry features. Our project in this paper is to seek a vantage point from which we can achieve such understanding and thus position ourselves to begin the selection process. The project thus requires one presupposition and two disclaimers:

- **Presupposition**: We speak throughout of *the conditional* and its logic and semantics. One moral of the market saturation noted above is that there is no *one* conditional, but a vast array of them, different ones suited to different logical, linguistic, and philosophical perspectives. Nevertheless, we will suppose throughout that there is a *core notion* of conditionality which runs through all or most conditionals (we leave open that some constructions people have called conditionals, or which bear the linguistic trappings typical of conditionals, are not in fact conditionals). It is this core notion that we want to limn – the goal is to extract a picture of the central logic and features, with an eye to the possibility that from the core may spring a family of conditionals further inflected via the introduction of various optional inferential features. We will use the symbol $\Rightarrow$ as a representation of the generic conditional we are considering.

- **First Disclaimer**: We aren't going to tell you which conditional to purchase, or even endorse any particular brands. No novel theory of the conditional is given in this paper, nor are any existing theories rejected. We think some morals for both old and new theories of the conditional are suggested by some of what we say, but the immediate goal is humble – we want to understand what we would be doing in selecting a conditional. We think that we, at least, have thus far too flimsy a grasp on the nature of the selection task to try to perform that task.

- **Second Disclaimer**: There is no novel linguistic data given in this paper (well, there are a couple of examples that we think are new, but they're not anything deeply significant), and there is no attempt to work up anything like a full semantic treatment of the conditional in a broader linguistic setting (in particular, we completely disregard the influential "ifs as restrictors" perspective).

We want, then, a more abstract perspective from which to (at least begin to) select a conditional. In order to reach this more abstract perspective, we start with a particular picture about the underlying nature of conditionals, and from this very general picture attempt to extract both some morals for the semantics of conditionals and some general cartographic orientation regarding the logical options. Our chosen picture is that most famously found in the "Ramsey test" conception of the conditional:

> If two people are arguing 'If p, will q?' and are both in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis about q. (Ramsey (1931), 247)

The Ramsey conception of the conditional is, of course, not a neutral starting point, and we are sensitive to the possibility that important perspectives on the conditional are immediately foreclosed by this opening. But it is hard to see how a view from nowhere is available here, and our particular development of the Ramsey conception will, we hope, show it to be a more generic and flexible launching point than one might initially think.

Any detailed implementation of the Ramsey conception narrowly construed must, of course, eventually muddy its hands with the details of what is involved in "adding [the antecedent] hypothetically to [a] stock of knowledge" and in "arguing on that basis about [the consequent]". Some aspects of the narrow construal, in fact, seem clearly misguided – surely there is no reason why it is only that which is *known* (rather than, say, believed) by an agent that is relevant (*qua* conditional-truth-evaluation) for his use of the antecedent in arguing for the consequent. Rather than plunging straight into the mud, we begin by attempting to extract from the Ramsey conception a minimal framework in which to situate further discussion. At a high level of generality, the Ramsey conception intends to trace the role of the conditional in reasoning, broadly construed to cover all normatively-governed modes of cognitive state change. Thought of in this manner, (Ramsey[2]) conditionals naturally conform to a rule of *conditional introduction*, or ⇒I. Such a rule allows one to reason to the truth of a conditional $A \Rightarrow B$ by starting an A-initiated path of reasoning terminating in B. At the moment, of course, "rule" is a bit of an overstatement – we have a general picture of the shape of an inferential procedure, but nothing so determinate that it deserves to be called a rule. Two further thoughts on ⇒I are immediately tempting:

1. One might think that ⇒I could be made fully determinate in the following manner. In a natural deduction system, allow a mode of *conditional subproof* which begins with a supplementary assumption of the desired antecedent A. Within the subproof, the full range of natural deduction proof rules can be deployed (either on the supplementary assumption or on other material introduced (e.g.) through further subproofs). Should this deployment eventuate in a claim B, then the subproof can be ended and a conditional $A \Rightarrow B$ can be inferred. This more specific version of ⇒I was endorsed as (half of) a characterization of the essence of the conditional in Bonevac et al. (2006). From this starting point, one could then build a family of conditionals in familiar fashion by controlling the permeability of the subproof barrier to outside information – complete impermeability would produce a minimal logic for the conditional on which the only available result was the theoremhood of $A \Rightarrow B$ whenever $A \vdash B$, permeability to all and only modal information would produce the strict conditional, permeability to information conditionally equivalent to the antecedent would produce the Lewis counterfactual, and full permeability would produce the material conditional (to list some options among many). However, in the current more abstract context, we will resist this particular conception of ⇒I. First, we do not want to suppose prematurely that the notion of normatively-governed cognitive state change that we are taking to underwrite the Ramsey conception must match up with the notion of deductive inference (perhaps some more restrictive mode of reasoning (giving rise, for example, to styles of relevance conditionals) or some less restrictive mode of reasoning (giving rise, for example, to common-sense reasoning nonmonotonic conditionals) should be deployed within the conditional subproof, or perhaps some

[2] We henceforth omit the qualification.

mode of reasoning orthogonal to deductive inference in strength should be deployed). Second, the eventual goal is to examine the source of various logical features of the conditional, and given this goal, to begin by underwriting the conditional with the full extent of deductive inference is to move too quickly and allow too many inferential features to emerge full-grown from the head of Zeus.

2. As a case in point of the above line of thought, one might suspect that *if absolutely nothing else*, the conditional introduction thought would lend pride of place to conditionals of the form $A \Rightarrow A$, as ones which ought to be endorsed even if we allow variation in the mode of normatively-governed transition from antecedent to consequent, since in this special case no (non-trivial) transition is even needed. Thus one thought that seems falls out of the Ramsey conception, as we are here limning it, is that a principle of *idempotence*, which identifies as theorems (or, weaker, as merely *true*) all conditionals of the form $A \Rightarrow A$, stands at the very center of our grasp of conditionality. However, one of our goals throughout this paper will be to emphasize that, on more careful examination, idempotence turns out to be much less centrally positioned in the logic of the conditional than one might originally have expected.

Conditional introduction does not, however, exhaust the commitments which intuitively flow from the Ramsey conception. If a conditional *expresses* (broadly) the availability of the consequent under the hypothesis of the antecedent, then, given the truth of a conditional, what it expresses should be the case, and hence the consequent should indeed be available under the hypothesis of the antecedent. Thus a conditional should also support a rule of *conditional elimination*, or $\Rightarrow$E, which allows that, when one is reasoning within a conditional subproof with initiating supplementary hypothesis A, and one has available (where what counts as availability is still up for debate) a conditional of the form $A \Rightarrow B$, then B can be introduced into the reasoning of the subproof. Were we to endorse the more particular reading, bruited above, of $\Rightarrow$I as backed by a notion of conditional subproof in which the normal inferential rules are made available for deployment on the antecedent, the result would be a concrete minimal conditional logic fully characterized by the axiom $A \Rightarrow A$ and the inferential rule that, given $A \Rightarrow B$ and C such that $A, B \vdash C$, one can infer $A \Rightarrow C$. (We would, for example, have a rule of *weakening of the consequent*, but would not have any rule (such as *transitivity*) which allowed transitions between conditionals with distinct antecedents).

- **Question**: Why endorse this rather constrained version of conditional elimination, rather than a more straightforward and more powerful rule of modus ponens, which allows the inferential transition from $A \rightarrow B$ and A to B? (Note that modus ponens entails the above rule of $\Rightarrow$E on

> the assumption that available conditionals (in whatever the final sense of availability turns out to be) can be reiterated into conditional subproofs whose supplementary hypothesis matches the conditional's antecedent)? An endorsement of modus ponens amounts, in essence, to a commitment to the available modes of inference being the same inside and outside the conditional subproof. Again, if one has bought into the deductive conception of the normatively-governed transitions underwriting the Ramsey conception, such a commitment may seem unexceptionable, but we want at this stage to avoid such a commitment, so modus ponens outruns the inferential commitments we take to be tacit in the Ramsey conception. Alternatively, it is well-known that the Lewis counterfactual has variants which do not produce modus ponens – those which lack the structural constraint of weak centering (the requirement that any world w be a member of the minimal sphere in its own sphere system). Stepping back slightly from the details of the Lewis semantics, one can think of the role of the antecedent in a conditional to be the introduction of a certain class of possibilities, a class which then plays some central role in the evaluation of the consequent. Thought of in this way, the rule of modus ponens is a reasonable one only if the actual world (when an A world) is guaranteed to be one of the class of possibilities introduced by an A antecedent. But absent a more detailed story of this selection function of the antecedent, such a conclusion is premature. (Note that on this way of thinking about things, we need not posit any mismatch between the inferential modes available inside and outside the conditional proof. Instead – if we think of propositions as sets of possible worlds – we see that the truth of A outside a conditional subproof and the appearance of A inside a conditional subproof will typically implicate different propositions (the set of *all* A worlds in the first case, but only of a class of A-introduced worlds in the second case), and propositions that, absent weak centering, need not be themselves inferentially related in a way that allows the cogent reasoning inside the conditional subproof from the one proposition to be carried over, even with identical notions of inferential connection, to cogent reasoning outside the subproof from different information.

Taken together, the ⇒I and ⇒E rules which we have argued fall out of the Ramsey conception of the conditional serve to bring that conception into engagement with another influential picture of the conditional introduced by Gentzen:

> To every logical symbol . . . belongs precisely one inference figure which "introduces" the symbol – as the terminal symbol of a formula – and one which "eliminates" it. . . . The introductions represent, as it were, the "definitions" of the symbols concerned, and the eliminations are no

> more, in the final analysis, than the consequences of these definitions. ... We were able to introduce the formula $\mathfrak{A} \supset \mathfrak{B}$ when there existed a derivation of $\mathfrak{B}$ from the assumption formula $\mathfrak{A}$. (Gentzen (1969), 80)

The Ramsey conception thus motivates a pair of (loosely conceived) inferential rules which characterize the conditional, in Gentzen style, via its introduction and elimination conditions. Of course, this minimal logic for the conditional leaves unanswered numerous questions about conditional inferences. Does the conditional support modus ponens, or modus tollens, or transitivity, or contraposition, or strengthening of the antecedent (the list, of course, can go on)? One approach at this point is to think of the minimal conditional logic as the foundation on which a richer logic can be built by selecting various options off of an inferential menu (the obvious analogy is with the superstructure of normal modal logics resting on **K**, although it remains to be seen whether any master semantic framework as deeply explanatory as the choice of features of the accessibility relation can be given in the conditional case). However, we prefer instead to begin by seeing whether these incidental inference rules can be extracted as consequences of the method of implementing the core Ramsey conception (one direction in which such an extraction might go has already been sketched above, in the idea of varying permeabilities to the conditional subproof barrier). We think there is a surprising amount of life to be had in fitting a broader range of inferential options into the core conception than might initially be expected, and think also that where inference rules resist integration into the core conception, there is the potential for learning more about the conditional (learning either that some inference rules, including perhaps inference rules that one pre-theoretically took to be paradigmatically conditional, are not in fact so deeply linked to the nature of conditionality (possibly even so weakly linked as to be wholly separate from), or alternatively that the Ramsey conception alone is insufficient to characterize the full notion of conditionality).

2. Antecedents

The Ramsey conception suggests that a semantic theory of the conditional needs to perform two tasks. First, some information state (to put things in the broadest possible terms) needs to be picked out. It is natural to think of this first task as being controlled by the antecedent of the conditional, and hence conforming to the Ramseyian thought of adding the antecedent to a prior stock of knowledge. Second, the selected information state must be deployed in testing (or being tested against) some further claim. It is natural to think of this second task as being controlled by the consequent of the conditional, and hence as conforming to the Ramseyian thought that the conditional involves "arguing about" the consequent. This correspondence between the Ramseyian division of labour and the syntactic division between antecedent and consequent is not a conceptual inevitability, of course, but it seems a reasonable working hypothesis.

2.1. Selection Functions

We thus begin with a discussion of the mechanisms through which the antecedent of the conditional plays the first Ramseyian role. There is a standard story regarding these mechanisms in place now, due to the work of Stalnaker and Lewis on *variably strict* conditionals, so our discussion on this half of the full Ramseyian package will be oriented toward properly contextualizing this story, rather than introducing substantive new ideas.

The heart of the Stalnaker/Lewis conditional semantics is the introduction of a *selection function*. The selection function is a mapping from a pair of a world and a proposition (which, given the background semantic setting, can be thought of as a set of worlds) to a proposition(/set of worlds). Given a selection function $f : W \times \wp(W) \mapsto \wp(W)$, we then have:

- $A \Rightarrow B$ is true at w iff $f(w, [A]) \subseteq [B]$ (where $[A]$ is the set of worlds at which A is true).

Historically, the selection function is best thought of as a maximally general answer to Goodman's problem (explored thoroughly in Goodman (1954)) of how we should properly amplify the antecedent of a (paradigmatically counterfactual) conditional. We want to secure the truth of:

- If I were to strike this match, it would light.

but we also know that there are match-striking circumstances (for example, ones in which the match is in a vacuum) that do not lead to lighting. Thus some facts about the actual situation (for example, that the match is in an oxygen-rich atmosphere) need to be preserved. However, other facts (for example, that the match is at a uniformly low temperature, or that it does not light) must not be preserved. We need some mechanism for transitioning to the overt (linguistically presented) content of the antecedent to a richer content that will allow proper evaluation of the conditional. (Note that this way of framing the issue implicitly takes the conditional to be, in some deeper sense, basically strict. The thought that possibilities in which the match is struck and does not light threaten the truth of the conditional, and hence need to be placed outside the bounds of consideration via an amplification of the antecedent (thereby creating a more restrictive sieve on worlds) is the thought that the conditional seeks universal covariation between antecedent and consequent. The Stalnaker/Lewis semantics in name abandons the strict for the variably strict, by no longer requiring the antecedent to carry the full burden of eliminating all possible counterexamples, but it remains close in spirit, by still requiring that the counterexamples be eliminated, and by requiring the antecedent at one remove to perform the filtering, via the intervening mechanism of the selection function.)

The difficulty in answering Goodman's problem has always been working out what stays and what goes when we entertain a counterfactual. In classic logician's style, the imposition of a selection function deals with this question by declining to answer it. Once a selection function is added to the semantics, one part of the specification of a model is a particular choice of selection function. The question of which selection function is the *right* selection function thus becomes a particular instance of the more general question of which model is right right model. But logicians are accustomed to deferring *this* question – determining the appropriate intended model is no part of logic proper, but is a task governed by the full range of capacities used to determine truth. What matters from the point of view of the logic is that the correct inferences are captured, and the capturing of inferences depends on the range of models, and not in any special way on the choice of intended model. Formally, the effect of the selection function is to wholly randomize the content of the antecedent, so that any inference involving conditionals of distinct antecedents fails. Thus with the addition of a selection function, it is easy to show that:

1. $A \Rightarrow B \nvDash (A \wedge C) \Rightarrow B$ (Consider a counterexample model $\mathcal{M}$ containing world w such that $f(w, [A]^{\mathcal{M}}) = [B]^{\mathcal{M}}$ and $f(w, [A]^{\mathcal{M}} \cap [C]^{\mathcal{M}}) = W_{\mathcal{M}} - [B]^{\mathcal{M}}$. So long as the valuation of $\mathcal{M}$ does not give $[C]^{\mathcal{M}} \subseteq [A]^{\mathcal{M}}$, such a model is always available.[3]

2. $A \rightarrow B,\ B \Rightarrow C \nvDash A \rightarrow C$ (Consider $\mathcal{M}$ and w such that $f(w, [A]^{\mathcal{M}}) = [B]^{\mathcal{M}} - [C]^{\mathcal{M}}$ and $f(w, [B]^{\mathcal{M}}) = [C]^{\mathcal{M}}$. So long as the valuation of $\mathcal{M}$ does not give $[B]^{\mathcal{M}} \subseteq [C]^{\mathcal{M}}$, such a model is always available.

The resulting logic is quite weak, and is (in its distinctively conditional fragment) fully axiomatized by the following principle:

- **Weakening of the Consequent**: From $A \Rightarrow B$ and $A \Rightarrow C$, infer $A \Rightarrow D$, where $B, C \vdash D$.[4]

Call this the **Minimal S-L Conditional Logic**. Of course, because the selection function introduces such a potent randomization into the semantics,

[3] Note that, as a result, the Stalnaker-Lewis semantics *will* support strengthening of the antecedent when the antecedent is strengthened (or "strengthened") by addition of a conjunct metaphysically necessitated by the original antecedent.

[4] A straightforward Henkin construction then shows completeness.

and thus yields such a minimal logic, it also threatens to produce a very unsatisfactory semantics for the conditional. Not *every* argument that involves variation in conditional antecedent appears automatically invalid, but no such argument can be captured by the pure selection function semantics. In the Lewis semantics, this illness is treated by imposed *ex nihilo* a highly structured arrangement of possibilities. Possible worlds are arranged into a *system of spheres* – a nested sequence of sets of worlds which collectively exhaust the range of possibilities (and are closed under union and non-empty intersection). Each world receives its own system of spheres. From these spheres, a selection function can immediately be recovered – we define $f(w, [A])$ to be the intersection of $[A]$ with the smallest sphere containing A worlds.[5] A selection function based on a system of spheres automatically and inevitably imposes certain coordinations between distinct antecedents. For example, such a selection function must endorse the inference:

- $A \Rightarrow B,\ B \Rightarrow A,\ A \Rightarrow C \vDash B \Rightarrow C$

At a given world w, $f(w, [A])$ picks out all A worlds in the smallest A-world containing sphere (henceforth, 'smallest A sphere') in w's sphere system. Given $w \vDash A \Rightarrow B$, all the A worlds in the smallest A sphere are B worlds. Also, $f(w, [B])$ picks out all the B worlds in the smallest B sphere, and given $w \vDash B \Rightarrow A$, all those B worlds are A worlds. But then the smallest A sphere must be the smallest B sphere – the smallest A sphere contains B worlds, and hence cannot be larger than the smallest B sphere, and by symmetric reasoning, the smallest B sphere cannot be larger than the smallest A sphere, so the two are identical. But within this smallest A, B sphere, all A worlds are B worlds, and all B worlds are A worlds, so $f(w, [A]) = f(w, [B])$. But, given the selection function semantics, this entails that A and B can be intersubstituted *salva veritate* in antecedent position, so from $A \Rightarrow C$, we can infer $B \Rightarrow C$.

This is only one example, but it is also a paradigmatic example – the general trend is for certain inferences involving distinct antecedents to become valid because the greatly constrained space of selection functions generated off of spheres require, given sufficient supplementary information, that certain antecedents coincide in their selection behaviour, making them appropriately intersubstitutable and thus allowing inferences that bridge across distinct an-

[5] We presuppose the Limit Assumption here: for every antecedent, there is a smallest sphere containing antecedent-supporting worlds. In fact, we will presuppose the stronger assumption that every system of spheres produces a nested sequence well-ordered by $\subseteq$. Under this stronger assumption, the requirement that a system of spheres be closed under non-empty intersections becomes superfluous. As the discussion below will show, starting from a selection function perspective makes the Limit Assumption all but inevitable. What to make, then, of the extra semantic possibility opened up by non-Limit-Assumption-supporting sphere systems is an interesting question deserving of further investigation.

tecedents. A natural deduction system can be designed for a robust Stalnaker-Lewis-style conditional by treating the conditional as governed by a rule of ⇒I in which external conditional information can be imported (via ⇒E) into the conditional subproof so long as the external conditional has an antecedent which is *conditionally equivalent* to the subproof's governing assumption (in the sense that the proof has already produced a conditional from external antecedent to subproof assumption and a conditional from subproof assumption to external antecedent). Call the resulting logic the **Pure S-L Conditional Logic**. This logic still falls short of Lewis's minimal system **V** in two ways:

1. We cannot derive certain interactions between the "would" conditional, the "might" conditional, and modalities. In particular, we cannot derive $A \,\Box\!\!\rightarrow B, A \,\Box\!\!\rightarrow \neg B \vDash \Box\neg A$ or the definitionally equivalent $A \,\Box\!\!\rightarrow B, \Diamond A \vDash A \,\Diamond\!\!\rightarrow B$.

2. We cannot derive a transitivity inference between "might" conditionals of the form $(A \vee B) \,\Diamond\!\!\rightarrow A,\ (B \vee C) \,\Diamond\!\!\rightarrow C \vDash (A \vee C) \,\Diamond\!\!\rightarrow A$.

It is tempting, nevertheless to think that the Pure S-L Conditional Logic isolates an interesting target of investigation, since these classes of additional features both crucially involve inferences that are not purely conditional (there are questions here about the sense in which the "might" conditional is a *conditional*, in the sense under investigation. These questions will be reopened below when we come to consideration of the Ramsey-style role of the consequent). However, the Lewis sphere system delivers not just the Pure S-L Conditional Logic, but invariably all of **V** – if the modals are taken to be the outer modals, in Lewis's sense, then both of the above inferences fall out of the sphere system. [6]

The Lewisian sphere system semantics thus provides a richer and more satisfying logic for the conditional, but as a *conceptual foundation* for the conditional, it is very unsatisfying. The greater inferential power comes in a single unarticulated package – we are given no insight either into how the individual components of the inferential strength are secured (this is the case for the central inferential core of the Lewis conditional, but not for some peripheral bits – modus ponens, for example, receives an independent foundation via the imposition of the weak centering requirement on the sphere systems) or into the reason that we should prefer this model-theoretic structure over some other. (Why, for example, a single linear ordering of spheres, rather than a branched system with two linear orderings sharing only a common root at

[6] Depending on which flavour of Lewis conditional is desired, some additional bells and whistles may be needed in the proof system, such as a rule of modus ponens (for weak centering), or the rule $A, B \vdash A \Rightarrow B$ (for strong centering), or the rule $A \Rightarrow B \vdash \Box(A \Rightarrow B)$ (for uniformity).

the minimal sphere?) The justification for the sphere-based semantics is (at this point) entirely consequentialist – we use this semantics because it validates the arguments we think should be validated, and refutes the arguments we think should be refuted. But if we are attempting to find a perspective from which we can understand *why* we ought to endorse or reject inferential forms, rather than just conforming to intuitive judgements of validity, this consequentialist approach is of no help.

2.2. Similarity and Spheres

Fortunately, Lewis offers reformulations of the sphere-based semantics which shed more conceptual light. One reformulation extracts the spheres out of an underlying comparison relation between words, typically glossed as a comparative similarity relation – we have $u \leq_w v$ to mean "u is more similar to w than v is". If we assume that $\leq_w$ is, for any w, a complete preordering (transitive and strongly connected), then we can extract a sphere system out of it.[7] Given any world v, define $S^w_v = \{u : u \leq_w v\}$. By transitivity, if $u \leq_w v$, then $S^w_u \subseteq S^w_v$, so by strong connectedness, the set of S^w_u form a nested sequence of sets. Closing this sequence under arbitrary unions and non-empty intersections yields a Lewisian sphere system for each world. So the Stalnaker-Lewis variably strict conditional can be thought of as underwritten by a notion of similarity, which serves to organize modal space in a way that allows for inferential interaction between distinct conditional antecedents. The similarity foundation is, we think, a conceptual advance over the direct imposition of the sphere system. However, we see two serious shortcomings of it:

1. The derivation of the sphere systems from the underlying similarity relation – and hence the production of the specific logic of the Stalnaker-Lewis condition – depends on a very particular, and insufficiently justified, level of abstraction. Lewis chooses to preserve, in the model theory, specifically the transitivity and strong connectedness of similarity. But why these features in particular? There are two linked problems. First, it is not obvious that similarity *has* these features. Strong connectedness is surely immediately suspect – we can certainly imagine there being two

[7] Lewis adds additional assumptions (i) guaranteeing that the base world of comparison w is always uniquely minimal in the similarity ranking (thereby enforcing strong centering in the induced sphere system and (ii) recognizing the possibility of an accessibility boundary beyond which worlds are not within the base world w's sphere system, and properly structuring the similarity relation to accommodate this boundary. We will continue to treat centeredness as a disjoint issue from the basic sphere system, and thus suppress this issue here. The possibility of an accessibility boundary matters only if we do not adopt Lewis's optional constraint of uniformity. We tacitly assume uniformity throughout, so – with the exception of one brief discussion in a subsequent footnote – we also suppress this aspect of Lewis's similarity relation.

objects *a* and *b* both of which are apt for comparison to a target object *c*, but which are so radically different from one another that they do not allow for direct comparison. (The number 3 may bear a certain similarity relation to my current thought, by virtue of that thought being about 3, and a certain neurological structure may bear a certain similarity relation to my current thought, by virtue of being the physical substrate of that thought, but it can nevertheless be that the number 3 and the neurological structure are themselves so radically and categorically different that the question of which is *more similar* to my thought cannot be sensibly addressed.[8]) Transitivity is also, although perhaps less obviously, suspect. If similarity is an attempt to summarize into a single dimension points of comparison along many dimensions – and surely it is – then one must worry about whether a similarity ranking between *a* and *b* based heavily on one subset of those many dimensions and another similarity ranking between *b* and *c* based heavily on a second disjoint subset of those many dimensions must inevitably translate into any particular similarity ranking between *a* and *c*, as transitivity would require. More abstractly, a multiple-feature-amalgamation model of similarity surely raises shades of Arrow's Theorem, and once Arrow's Theorem is invoked, the denial of transitivity looms large as one potential member of an inconsistent set of constraints to be rejected.

Second, even if comparative similarity *is* transitive and strongly connected, it is unclear why these two particular features of the comparative similarity relation are privileged into the model theory. Similarity, for example, is plausibly context-sensitive (in some complex and ill-understood way) – why do we not also write context-sensitivity into the abstract structural skeleton of the similarity relation? More troublingly: Lewis's work on conditionals leads to something of a cottage industry setting out a more detailed theory of similarity capable of producing intuitively acceptable truth conditions for problematic conditionals (see Fine (1975) and Bennett (1974) for early examples of this industry, and Bennett (2003), chapters 11-15 for a recent discussion of the industry state of the art). Thus, for example, we are told that (in the relevant-to-conditionals sense of similarity) similarity of physical law trumps similarity of contingent fact. Should we require that this fact be written into

[8] We can, of course, also imagine that there are objects so outré that they are not even apt for similarity comparison to *c*. Lewis does not directly have the formal means to accommodate such outliers, due to his requirement of strong connectedness, but he can simulate a treatment by placing outlier worlds outside the accessibility boundary, and treating all such worlds as uniformly maximally dissimilar to the base world *w*. However, in doing so, Lewis must assume that two notions coincide: (i) being a world so outré that it is not apt for similarity comparison to *w*, and (ii) being a world which never features in counterfactual reasoning from *w*. It is unclear what justifies the requisite coincidence assumption.

the model theory? We are also told that similarities in the temporal past of the antecedent weigh more heavily that similarities in the temporal future of the antecedent – again, this constraint could be written into the model-theoretic treatment of the similarity relation. As such additional constraints are added to the similarity relation, the resulting organization of worlds will become correspondingly more structured, and additional inferential features of the conditional may emerge. In the limit, the danger emerges that the similarity relation can be so thoroughly characterized that the idea of a merely structural skeleton of the relation built into the model theory, allowing for model-by-model variation in realization and hence a substantive question regarding the selection of an intended model may be made moot, as only one organization of the worlds into spheres is compatible with the enriched structuring similarity relation.

2. The use of a similarity relation to build the sphere system, and hence to secure the desired coordination between distinct antecedents, ends up effectively displacing the task of determining a body of information against which the consequent is to be tested (that is, performing the first half of the Ramsey procedure) away from the antecedent and onto the static similarity structure. This is not to say that the antecedent plays no role in determining the testing information – of course, it must, if there is to be truth-conditional variation between conditionals with identical consequents and distinct antecedents. But the similarity relation, which does all the work of generating the sphere system, is a relation directly between worlds, and a relation which in itself takes no account of the antecedent. So the determination of a body of information proceeds by using the antecedent as a tool to select a region in a topography that is itself based entirely on direct comparisons between worlds, and is held constant across all choices of antecedent. It is not perfectly clear how sharp a worry exists here, but the general shape of the concern is two-fold. First, by extracting the needed body of information off of an antecedent-independent structure, we run the risk of making less conceptually clear why we get the antecedent-relating inferential patterns we do out of the Stalnaker-Lewis conditional. Suppose we want to know *why* two conditionally equivalent antecedents (say, A and B) always select the same body of information. The answer crucially involves the positioning of the A and B worlds in the sphere system, which in turn involves the comparative similarity relations between A and B worlds. But those similarity relations are themselves directly world-generated, and not in any way inflected by the specific propositions in question, so we may end up lacking sufficient insight into why *these particular antecedents* intersubstitute in the way they do in conditional contexts. Second, by downplaying the role of the antecedent – in particular, by having the antecedent play only

a secondary role of picking from a pre-established structure – the final picture diverges somewhat from the desired Ramseyian conception, on which the antecedent is playing a primary and active role in the shaping of the tested body of information.

The rather nebulous point can perhaps be given a bit more bite via the following problem:

> **Dragging Problem**: Consider conditionals with antecedents which are realized only in worlds very dissimilar to ours. If the semantic role of the antecedent is achieved via a comparative similarity relation, then the evaluation of such conditionals (relative to our world) will bring into consideration a very wide range of worlds. The result is that it will be relatively difficult for such conditionals to be true, since the consequent will have to hold in a very wide range of worlds. In some cases, this is a desirable effect – a conditional such as:
>
> - If Ralph Nader had won the 2000 presidential election, then the Dow would currently be above 13,000.
>
> ought to be at best highly suspect, since in the outlandish worlds containing a Nader victory, much may be up for grabs. But this "dragging" effect looks undesirable when the outlandishness of the antecedent seems irrelevant to the consequent. Thus consider:
>
> - If there were a region of five-dimensional space in the Andromeda galaxy, then it would be cold in Michigan in the winter.
>
> This looks, we think, true. But the dragging effect ought also to defeat it, since we must extend to quite odd worlds to get the relevant region of five-dimensional space, and surely within the very large sphere thus delimited, there will be worlds in which Michigan is warm in the winter (not because of the five-dimensional space, but because of (for example) a much less odd slight increase in energy output from the sun). The dragging effects show up even more dramatically with dynamic systems in which entertainment of antecedents affects the context for downstream conditionals – in such systems, if we utter "If there were a region of five-dimensional space in the Andromeda galaxy, then that galaxy would be quite odd", we could not proceed truthfully to utter "If it were snow in Michigan in the winter, the snow would be white".

The dragging problem is a manifestation of the univocality and antecedent-insensitivity of the similarity measure – since the similarity structure is in place independently of antecedent, all the antecedent can do is pick a location in the similarity structure, and in doing so, it will typically drag into consideration possibilities that are intuitively irrelevant to consideration of the consequent. One might compare the situation here to that found under the Asher-Morreau normality conditional, in which the antecedent selects via an antecedent-specific similarity relation (roughly, in consideration of $A \Rightarrow B$, one considers all maximally A-normal worlds, where A-normality is not just the intersection of the presence of A with a background pure normality ordering. Thus consideration of situations highly abnormal in one respect need not induce any disruption of regularities in other respects).

2.3. Spheres and Selection Functions

A sphere system can also be extracted out of a selection function. Not just any selection function, of course – given the full range of selection functions, the minimal conditional logic results, while a sphere system enforces a much stronger logic. But with adequate constraints, a selection function will suffice to generate a sphere system. Lewis lists the following constraints:

1. If $w \vDash A$, then $f(w, [A]) = \{w\}$
2. If $[A] \subseteq [B]$ and $f(w, [A]) \neq \emptyset$, then $f(w, [B]) \neq \emptyset$
3. $f(w, [A]) \subseteq [A]$
4. If $[A] \subset [B]$ and $[A] \cap f(w, [B]) \neq \emptyset$, then $f(w, [A]) = [A] \cap f(w, [B])$

The first constraint enforces strong centering of the sphere system – as usual, we will simply bracket discussion of this constraint (it is worth noting that strong centering needs its own independent parameter on the selection function to produce, and does not in any natural way fall out of the mere production of a sphere system).[9] The role of the second constraint is less straightforward. A stronger version of the second constraint:

- If $[A] \neq \emptyset$, then $f(w, [A]) \neq \emptyset$

imposes the Lewisian constraint of absoluteness, by requiring that each world occur in each sphere system.[10] The weaker version does not require absoluteness. Absent absoluteness, Lewis allows that there be a "accessibility

[9] If we want weak centering instead of strong centering, this constraint should be weakened to:

- If $w \vDash A$, then $w \in f(w, [A])$

[10] Almost. In fact, this strengthened version of the second constraint is compatible with the existence of "elusive worlds" – worlds such that any characterization of them is also satisfied by some closer world, thus preventing them from ever being selected. (Of course, such elusive worlds are possible only given that the language of selection is $L_{\omega\omega}$, rather than $L_{\omega\omega_1}$.)

boundary" beyond which worlds never enter into the entertainment of conditional antecedents. In a sphere system, no further axiomatic work need be done to enforce the accessibility boundary – certain worlds simply are not placed in the sphere system, and thus, via the conditional semantics, can never be accessed by any antecedent. But when we operate directly with a selection function, there is substantive work to be done by the selection function to create a coherent boundary. The notion of a boundary plays a role in the semantics only insofar as it makes unavailable the truth of certain claims. In order for a distribution of unavailabilities to have a properly "boundary-like" shape, there must be proper coordination between the unavailabilities. If A is unavailable, for example, $A \wedge B$ ought also to be unavailable. The second condition thus ensures that the selection function will generate a true boundary. Since the second condition is implicated, in essence, in forcing the selection function to treat input propositions as world-like, one might expect that it would also play the role of requiring that sets of worlds selected by inferentially related propositions be appropriately nested, but it turns out that this condition on its own does not suffice, since it does nothing to require nesting behaviour in selected propositions. Below we will see that the second and fourth conditions taken together do guarantee appropriate nesting, however.

The third condition enforces the natural-seeming requirement that the selection function select A worlds for the input A. (Such a requirement, for example, is surely mandated by the Goodman-problem conception of the role of the selection function.) This condition thus serves to guarantee the truth of idempotence conditionals $A \Rightarrow A$. As such, it would seem to be a hermetic requirement along the lines of the first requirement (enforcing centering), but we will see shortly that it also plays a crucial role in the integrative reasoning showing that a sphere system emerges from the selection function.

The fourth condition is, from a certain point of view, where all the action is. One way to put some substance to that claim: recall that the ultimate role of constraints on the selection function is to allow non-trivial inferential engagement between distinct antecedents, and that the central form of such engagement is the requirement that conditionally equivalent antecedents be intersubstitutable. That intersubstitution requirement amounts to the following constraint on the selection function:

- If $f(w, [A]) \subseteq [B]$ and $f(w, [B]) \subseteq [A]$, then $f(w, [A]) = f(w, [B])$.

But in fact this constraint follows from Lewis's fourth constraint (not alone, but with the fourth constraint doing the substantive work):

Claim: Let $f(\cdot, \cdot)$ be a Lewis selection function. Then if $f(w, [A]) \subseteq [B]$ and $f(w, [B]) \subseteq [A]$, then $f(w, [A]) = f(w, [B])$.

Proof: If $f(w, [A])$ and $f(w, [B])$ are both empty, the claim is trivial, so assume at least one is non-empty. Without loss of generality, take $f(w, [A]) \neq \emptyset$. By the third constraint, $f(w, [A]) \subseteq [A]$, so $f(w, [A]) \cap [A] \neq \emptyset$. But given $f(w, [A]) \subseteq [B]$, we have $f(w, [A]) = f(w, [A]) \cap [B]$, so $f(w, [A]) \cap [B] \cap [A] \neq \emptyset$. $[A] \cap [B] = [A \wedge B]$, so $[A \wedge B] \cap f(w, [A]) \neq \emptyset$. But since $[A \wedge B] \subseteq [A]$, by the fourth condition we have $f(w, [A \wedge B]) = [A \wedge B] \cap f(w, [A]) = [A] \cap [B] \cap f(w, [A]) = [A] \cap f(w, [A]) = f(w, [A])$.

Since $f(w, [A \wedge B]) \neq \emptyset$ and $[A \wedge B] \subset [B]$, by the second constraint we have $f(w, [B]) \neq \emptyset$. Hence $[A \wedge B] \cap f(w, [B]) \neq \emptyset$. But then, reasoning as above, we have by the fourth constraint $f(w, [A \wedge B]) = [A \wedge B] \cap f(w, [B]) = [A] \cap [B] \cap f(w, [B]) = [A] \cap f(w, [B]) = f(w, [B])$. Combining these equalities, we have $f(w, [A]) = f(w, [A \wedge B]) = f(w, [B])$ as desired.

The fourth condition then requires that when propositions are appropriately related, the propositions selected by them are appropriately nested. The twin appropriateness conditions require some care to state – we cannot simply require that if $[A] \subseteq [B]$, then $f(w, [B]) \subseteq f(w, [A])$, because the selected B worlds might include worlds at which A, the logically stronger proposition, fails.[11] But we can get a result close to this: if we can split a logically weaker proposition exhaustively into two logically stronger propositions, then we can require that the worlds selected by the weaker proposition are a subset of the worlds selected by one of the two stronger propositions:

Lemma 1: If $f(\cdot, \cdot)$ is a Lewis selection function, then for any B and C, either $f(w, [B]) \subseteq f(w, [B \vee C])$ or $f(w, [C]) \subseteq f(w, [B \vee C])$.

Proof: If either $f(w, [B])$ or $f(w, [C])$ is empty, the claim is trivial, so assume neither is empty. Since $[B] \subseteq [B \vee C]$, then (by the second constraint) $f(w, [B \vee C])$ is also non-empty. Pick $u \in f(w, [B \vee C])$. By the third constraint $u \in [B \vee C]$ so either $u \in [B]$ or $u \in [C]$. Suppose $u \in [B]$. Then $[B] \cap f(w, [B \vee C]) \neq \emptyset$. Since also $[B] \subseteq [B \vee C]$, by the fourth constraint $f(w, [B]) = [B] \cap f(w, [B \vee C])$. Hence $f(w, [B]) \subseteq f(w, [B \vee C])$. Similarly, if $u \in [C]$, we have $f(w, [C]) \subseteq f(w, [B \vee C])$.

Once a selection function satisfies all four constraints and is a Lewis selection function, a system of spheres can be built off of it. Following Lewis, we take a sphere (in the system around w) to be any set S satisfying the following two constraints:

[11] We of course do not want $[A] \subseteq [B]$ to yield $f(w, [A]) \subseteq f(w, [B])$, since this would amount to a rule of strengthening of the antecedent.

1. For all $u \in S$, there is a sentence A such that $u \in f_w(A)$.[12]
2. For all sentences A, if $[A] \cap S \neq \emptyset$, then $f_w(A) \subseteq S$.

It remains to show that "spheres" defined in this sense form a proper sphere system. Setting aside the closure conditions, and allowing centering to be directly handled by the first constraint, the only issue is whether spheres nest properly – whether, given any spheres S_1 and S_2 (induced by $f(w, \cdot)$), we have $S_1 \subseteq S_2$ or $S_2 \subseteq S_1$. And, in fact, this result does follow:

Claim: If S_1 and S_2 are spheres induced by $f(w, \cdot)$, then either $S_1 \subseteq S_2$ or $S_2 \subseteq S_1$.

Proof: We begin by proving:

Lemma 2: Given any B and C, and S induced by $f(w, \cdot)$, if $S \cap f(w, [B]) \neq \emptyset$, then $f(w, [B \vee C]) \subseteq S$.

Proof: Take $u \in S \cap f(w, [B])$. Since $f(w, [B]) \subseteq [B]$, $u \in [B]$, and hence $u \in [B \vee C]$. So $S \cap [B \vee C] \neq \emptyset$, and by the second requirement on induced spheres, $f(w, [B \vee C] \subseteq S$.

Given this lemma, the desired result follows quickly. Suppose, for indirect proof, that $S_1 \not\subseteq S_2$ and $S_2 \not\subseteq S_1$. Then there is $u \in S_1$, $u \notin S_2$, and also $v \in S_2$, $v \notin S_1$. By the first requirement on induced spheres, there are then B and C such that $u \in f(w, [B])$ and $v \in f(w, [C])$. Since $u \in f(w, [B])$, $u \in [B]$, so $S_1 \cap [B] \neq \emptyset$, and $f(w, [B]) \subseteq S_1$. Similarly, reasoning from v, $f(w, [C]) \subseteq S_2$. By Lemma 2, $f(w, [B \vee C]) \subseteq S_1$, S_2, and by Lemma 1, either $f(w, [B]) \subseteq f(w, [B \vee C])$ or $f(w, [C]) \subseteq f(w, [B \vee C])$. Thus either $f(w, [B]) \subseteq S_2$, or $f(w, [C]) \subseteq S_1$. But then either $u \in S_2$ or $v \in S_1$, contradicting the selection of u and v.

So a Lewis selection function will generate the desired sphere system, and hence ground the desired inferential behaviour of the conditional (focusing, still, on the role of the antecedent). What remains unclear is whether some backing conception of the nature of a conditional, or of the role of the antecedent, can motivate the specific constraints that Lewis imposes on the selection function. Before turning to a proposed route of motivation, we first observe a second method of structuring the constraints on the selection function. As noted above, a Lewis selection function is sufficient to guarantee that $f(w, [A]) \subseteq [B]$ and $f(w, [B]) \subseteq [A]$, then $f(w, [A]) = f(w, [B])$. This feature of the selection function underwrites the following inference rule in a natural deduction system:

[12] Note that this condition amounts to a requirement that there are no "elusive worlds" in the sense of footnote 10. It would be nice if we had a clearer handle on whether it was good or bad that there be elusive worlds, especially – given that the selection function here is meant to encapsulate our grasp of the role of the antecedent in the function of the conditional – a handle which derived from our grasp of the conditional.

- In a conditional subproof with assumption for conditional proof of A, one can employ a rule of conditional elimination to introduce B into the subproof if, external to the subproof, one has a counterfactual of the form $C \Rightarrow B$, where C is shown in the proof to be conditionally equivalent to A.

Since this inference rule is the core of the Pure S-L Conditional Logic, it is then tempting to consider placing it *directly* as a constraint on the selection function, and then dropping at least the third and fourth Lewis constraints. (The first constraint will need to be retained if we want a weakly centered system, and the second constraint will be needed if we want our conditional to be idempotent.) However, the resulting class of selection functions are no longer sufficiently structured to guarantee the extractability of a sphere system. This is easily seen by taking the identity map as a selection function (for fixed w) in a space of more than one world.

One advantage that thus immediately drops out of using an axiomatically constrained selection function as a foundation for the Lewis-Stalnaker conditional, rather than immediately the sphere system, is that we can trace the incremental development of the inferential strength of the conditional. Call a selection function with Lewis conditions (1) and (2), along with if $f(w, [A]) \subseteq [B]$ and $f(w, [B]) \subseteq [A]$, then $f(w, [A]) = f(w, [B])$, a *weakly Lewis selection function*, and a selection function with Lewis conditions (1)-(4) a *strongly Lewis selection function*. Then a weakly Lewis selection function generates the Pure S-L Conditional Logic, while a strongly Lewis selection function generates the Lewis system **VC** (the minimal Lewis system **V** can easily be extracted by dropping the first selection function constraint). We thus have the potential for gaining an understanding of why conditionals have the inferential features they do, or tools for profitably disputing over particular inferential features. But the same worry that haunted the attempt to use comparative similarity as a conceptual foundation persists here – if anything, the worry is stronger. So far we have only a list of axiomatic constraints on a selection function, and no (non-consequentialist) account of why we should accept these constraints rather than others. The selection function can be conceptually grounded in the Goodman problem – as the tool for providing the needed amplification of the antecedent – but then, as with similarity, we face the question of why only certain structural features of what on the face of it is a highly structured notion get preserved into the semantics. (Why, for example, do we not encode into the semantics the thought that law-like claims are preserved preferentially with respect to matters of accidental fact?)

However, a different approach to extracting the needed structural features of a selection function shows more promise. Returning to the underlying Ramsey conception, we can think of the antecedent as guiding the alteration of an information state under the incoming influence of a new piece of data. One particular model for this process is then belief update models, although our maximally general perspective on the Ramsey conception means that we do not need to be bound to thinking of the process specifically as one of belief dynamics (we could, for example, make room for a distinction between the dynamics of belief and of hypothetical reasoning). The general shape of the proposal is this: suppose we have some operation $\bigoplus$ which updates information states with new data. Then $\bigoplus$ is a binary operation, and if we adopt the simplest model of treating information states as sets of worlds, and data as propositions modelled also as sets of worlds, then $\bigoplus$ maps from a pair of sets of worlds to a set of worlds. Formally, this is almost what we want from a selection function, and it yields the following natural semantics for conditionals:

- $A \Rightarrow B$ is true at information state $\mathcal{I}$ iff $\mathcal{I} \bigoplus [A] \subseteq [B]$. (Alternatively, we can think of information states as being sets of propositions, combined with an abuse of notation on which a sentence serves to name the proposition it expresses, and write $A \in \mathcal{I}$ to mean $[A] \subseteq \mathcal{I}$. Then we have $A \Rightarrow B$ is true at $\mathcal{I}$ iff $B \in \mathcal{I} \bigoplus A$.

To design the semantics for conditionals in this way is, of course, to make their truth relative to information states, rather than relative to worlds (as in the Lewis-Stalnaker semantics). But to make conditionals true relative to information states is then to have them sit poorly with the rest of the semantics, which uses truth relative to a world as the basic semantic coin. Four possible reactions to this difficulty:

1. **Quarantine**: One can simply accept that conditionals traffic in a different semantic coin from non-conditional claims. To avoid semantic type clashes, one then blocks semantic interaction between conditionals and other expressions. Such quarantining would, for example, block the negating of conditionals, or the embedding of conditionals in conditionals.

2. **Lift**: One can engage in currency exchange to make information states the basic semantic coin for the whole language. This can be done in a trivial way, by taking a claim A, for which we have a truth-at-a-world definition, and defining A to be true at information state $\mathcal{I}$ iff A is true at w for all $w \in \mathcal{I}$. Things then proceed in the usual way – $A \wedge B$, for example, is true at an information state if that information state is the intersection of two information states at which A and B are true, respectively. One has at this

point the option of using the general move to truth at information states as a tool for dropping the continuity of classical semantics – the requirement that the truth of a sentence at an information state supervene on the truth of the sentence at (the singleton of) each world in that information state.[13] The result would then be a dynamic treatment, of some flavour or other, of the conditional. Much of what we say here should carry over to such treatments.

3. **Lower**: One can exchange currency in the other direction by making worlds the basic semantic coin. The difficulty is then one of converting truth at an information state to truth at a world, which in turn requires associating worlds with information states (so that a conditional $A \Rightarrow B$ can be true at w just in case it is true at the information state $\mathcal{I}$ associated with w). Formally, the tool for achieving such association is the accessibility relation, but the limitation (for cardinality reasons) is that each world can be associated only with one information state, and one then wants a justification for associating a world with one information state rather than another.

4. **Dominate**: One can adopt the following principle regarding evaluation of conditionals:

    - If two people P_1 and P_2 are in information states $\mathcal{I}_1$ and $\mathcal{I}_2$ such that a conditional $A \Rightarrow B$ has different truth values in $\mathcal{I}_1$ and $\mathcal{I}_2$, and information state $\mathcal{I}_1$ is richer than information state $\mathcal{I}_2$ (so that $\mathcal{I}_1 \subset \mathcal{I}_2$), then $\mathcal{I}_1$'s assessment of the truth value of the conditional should be accepted rather than $\mathcal{I}_2$'s.

    This principle is deliberately vague on the nature of the norm and of acceptance, but it is meant to help explain a common reaction to Gibbard (1981)'s Sly Pete example: thinking that Jack's assessment of the conditional is to be preferred to Zack's, because Zack would drop his claim of truth if he were aware of the facts Jack knows, whereas Jack would persist in thinking the conditional false, even if he knew what Zack knows. One might then be pushed by this principle to adopt the claim that, although conditionals have a superficial dependence on information states, they are in fact ultimately governed only by maximally determinate information states – i.e., worlds. Formally, this amounts to the common technique of introducing a type lowering operation by having that operation replace a (higher-type) principle ultrafilter with the (lower type) generator of that ultrafilter.

[13] See van Benthem (1986) for details on continuity and its relation to the classical semantic picture.

We will henceforth bracket questions of how to deal with the type-mismatch induced by an information-state-based semantics for conditionals, and for the remainder of this section will proceed as if assigning truth to conditionals at information states was unproblematic.

The central question thus becomes: what is the nature of the $\bigoplus$ belief update operation? A natural starting point is the following:

- **Stubborn Update**: $\mathcal{I} \bigoplus A = \mathcal{I} \cap [A]$

The Stubborn Updater takes on board new information simply by adding it to his current information state. Stubborn Update more or less produces the strict conditional. The C.I. Lewis axioms for S3, for example, are all supported by Stubborn Update. (Caveats: in order to avoid the semantic coin problems discussed above, imbedded conditionals are read instead as inferences, and modals are implemented as true at an information state just in case they are true in the S5 modal model straightforwardly corresponding to the information state.) Stubborn Update yields Strengthening of the Antecedent – if $A \Rightarrow B$ is true in state $\mathcal{I}$, then $\mathcal{I} \bigoplus A = \mathcal{I} \cap [A] \subseteq [B]$. But then $(A \wedge C) \Rightarrow B$ requires that $\mathcal{I} \bigoplus (A \wedge C) = \mathcal{I} \cap [A \wedge C] = \mathcal{I} \cap [A] \cap [C] \subseteq \mathcal{I} \cap [A]$ be a subset of $[B]$, which follows immediately from the previous claim. Stubborn Update falls short of a standard strict conditional only in its inferential interactions with more robust modals which look invariably to the full range of worlds, rather than just to the worlds in the current information state – here Stubborn Update's insistence on looking at a smaller and smaller region of modal space will cause certain characteristic inferences to fail (for example, Stubborn Update does not support, with the robust modal, $A \Rightarrow B, \Diamond A \vDash \Diamond B$, since $A \Rightarrow B$ can hold by virtue of the local absence of A worlds in $\mathcal{I}$, leading Stubborn Update to test B against an empty information state – a test which will be passed whether B represents a genuine (robust) possibility or not). To get the pure strict conditional, one would need the rather more idiosyncratic:

- **Distractable Update**: $\mathcal{I} \bigoplus A = \mathcal{W} \cap [A]$, where $\mathcal{W}$ is the full set of worlds.

The distractable updater forgets everything he had learned when encountering new information A, and hence comes away believing only A.

Neither Stubborn nor (the henceforth ignored) Distractable Update are realistic pictures of belief dynamics, and the nature of their irrealism then sheds light on what a variably strict conditional needs to do to overcome the philosophical shortcomings of the strict conditional. The central problem with Stubborn Update is that it makes no allowances for nonmonotonicity in information states – there is no notion of removing old information in the face of contradicting new information. If one is attempting to model the use of a conditional to contemplate contrary-to-fact antecedents, this shortcoming will be highlighted, and will produce unacceptable inferential features. But suppose we try to hold on at least to a bit of stubbornness, with the principle:

- **Weak Conservativity**: If $A \in \mathcal{I}$, then $\mathcal{I} \bigoplus A = \mathcal{I}$.

The Weakly Conservative updater needn't always hold on to old information, but he at least tries to, in the following minimal sense: if he is asked to update his belief state with a piece of information that he already accepts, then he simply does nothing.

Weak Conservativity won't buy us much, but if it is combined with two principles regarding conjunctions, a nice result drops out. So consider:

- **Conjunctive Update**: $\mathcal{I} \bigoplus (A \wedge B) = (\mathcal{I} \bigoplus A) \bigoplus B$
- **Conjunctive Symmetry**: $\mathcal{I} \bigoplus (A \wedge B) = \mathcal{I} \bigoplus (B \wedge A)$

Conjuctive Update has it that the way to alter one's beliefs in the face of conjunctive information is to alter one's beliefs in the face of the first conjunct, and then to alter those altered beliefs in the face of the second conjunct. Conjunctive Symmetry then holds that the order of conjuncts makes no difference to the outcome of updating. This principle is redundant in the current framework, since the update operation acts on sets of worlds, and $(A \wedge B)$ and $(B \wedge A)$ pick out the same set of worlds. But in a more finegrained framework, it could be a substantive requirement, and we highlight it here because one natural path of development involves rejecting it.

We can now show the following:

Claim: Given Weak Conservatity, Conjunctive Update, and Conjunctive Symmetry, if $\mathcal{I} \bigoplus A \subseteq [B]$ and $\mathcal{I} \bigoplus B \subseteq [A]$, then $\mathcal{I} \bigoplus A = \mathcal{I} \bigoplus B$.

Proof: Since $\mathcal{I} \bigoplus A \subseteq [B]$, $B \in \mathcal{I} \bigoplus A$, and by Weak Conservativity, $\mathcal{I} \bigoplus A = (\mathcal{I} \bigoplus A) \bigoplus B$. By Conjunctive Update, $(\mathcal{I} \bigoplus A) \bigoplus B = \mathcal{I} \bigoplus (A \wedge B)$, which in turn by Conjunctive Symmetry is $\mathcal{I} \bigoplus (B \wedge A)$. Applying Conjunctive Update again yields $(\mathcal{I} \bigoplus B) \bigoplus A$. But since $\mathcal{I} \bigoplus B \subseteq [A]$, we have $A \in \mathcal{I} \bigoplus B$, and by Weak Conservativity, $(\mathcal{I} \bigoplus B) \bigoplus A = \mathcal{I} \bigoplus B$. Chaining together identities, we have $\mathcal{I} \bigoplus A = \mathcal{I} \bigoplus B$, as desired.

This result is clearly the belief-update analog of the third selection function constraint in a weakly Lewis selection function. Thus some natural constraints on a belief update operation yield a central requirement for a selection function to produce the Pure S-L Conditional Logic. It is worth noting, however, that Conjunctive Symmetry will look objectionable from a dynamic perspective, and that it is specifically from a dynamic perspective that Conjunctive Update looks most tempting. Thus it is not immediately clear that there is a single coherent position that fully motivates the constitutive components of the derivation.

The second constraint of both the weakly and strongly Lewis selection function – $f(w, [A]) \subseteq A$ – also finds a natural reading as a belief update constraint. In the belief update context, it becomes:

- **Success**: $A \in \mathcal{I} \bigoplus A$

Success can look irresistible, but there are both general and specific reasons for calling it into question:

1. **General**: A Harman-style (Harman (1986)) distinction between rules of implications and rules inference can help call Success into question. Of course from the information that *A*, *A* is implied. But it does not follow that, upon receipt of the information that *A*, one is rationally obligated to transition to a cognitive state that has it that *A*. One might, instead, have the rational option of rejecting the information that *A*. The option might even, in the right cognitive circumstances, become an obligation.

2. **Specific**: Dynamic treatments of Moore-paradoxical sentences, such as that in Gillies (2001), take these sentences to have the feature that a belief state that is updated with one of them nevertheless becomes a belief state which does not support that sentence. So one can coherently adopt "it might be that p and $\neg p$" into one's belief state, by first having one's prior belief state pass the "might p" test and then updating with $\neg p$, but the resulting output belief state will not itself any longer pass the "might p" test, and hence will not support the Moore-paradoxical input to update.

If Success is taken as the grounds for the principle $f(w, [A]) \subseteq [A]$, then it does all of the work of grounding the feature of idempotence for conditionals – the rest of the behaviour of the antecedent can proceed unaltered with or without Success, with idempotence then standing or falling respectively.

The first constraint of both the weakly and strongly Lewis selection function is the strong centering constraint that if $w \in [A]$, then $f(w, [A]) = \{w\}$, or (alternatively) the weak centering constraint that if $w \in [A]$, then $w \in f(w, [A])$. However, neither of these constraints finds a natural home in the belief update setting, in large part because of the truth-at-a-world/truth-at-a-state mismatch above – the antecedent *A* takes its truth value at a world, and this fact is central in the statement of the centering constraints on a selection function; as a result, no natural constraint on belief updating matches the centering constraint, because there is no analog in the belief update function to the privileged role of w. If we lift truth of non-conditionals up to state level, then the analog of the constraint seems to be that if $A \in \mathcal{I}$, then $A \in \mathcal{I} \bigoplus A$. (It is now unclear how to distinguish between weak and strong centering.) This constraint will follow from the combination of Success and Weak Conservativity.

So much for weakly Lewis selection functions; the remaining issue is now the two further constraints on a strongly Lewis selection function:

1. If $[A] \subseteq [B]$ and $f(w, [A]) \neq \emptyset$, then $f(w, [B]) \neq \emptyset$. (Or, if we want absoluteness, if $[A] \neq \emptyset$, then $f(w, [A]) \neq \emptyset$.)

2. If $[A] \subseteq [B]$ and $[A] \cap f(w, [B]) \neq \emptyset$, then $f(w, [A]) = [A] \cap f(w, [B])$

Consider first the stronger absoluteness constraint that if $[A] \neq \emptyset$, then $f(w, [A]) \neq \emptyset$. Translated into belief update talk, this amounts to a principle of:

- **Recovery**: If one update one's belief state with any non-contradictory piece of information, one always ends up in a coherent belief state.

Recovery is a plausible but not irresistible constraint. It gains bite when coupled with Success – absent Success, Recovery can more or less be secured by having a last-resort strategy of updating-by-rejecting if an incoming piece of information is too hard to fit together with a current belief scheme. In the context of Success, Recovery bears a certain resemblance to the presence of the **D** axiom in deontic logic – it amounts to the assumption that there are no irresolvable conflicts in doxastic commitments, so that one can always find a way forward that is doxastically permissible.

The weaker constraint that if $[A] \subseteq [B]$ and $f(w, [A]) \neq \emptyset$, then $f(w, [B]) \neq \emptyset$ follows less naturally from pure belief update constraints. It is easy enough to see how to translate it into belief update language:

- **Tolerance**: If $[A] \subseteq [B]$, and $\mathcal{I} \bigoplus A$ does not crash, then $\mathcal{I} \bigoplus B$ does not crash.

But why should we accept such a principle? One plausible line of thought is that we derive it from another conservativity principle:

- **Strong Conservativity**: If $\neg A \notin \mathcal{I}$, then $\mathcal{I} \bigoplus A \neq \emptyset$.

If, then, we acknowledge a claim as an open possibility (we have not already settled on its negation), then we can take that claim on board without crashing our belief state. Suppose $[A] \subseteq [B]$ and $\mathcal{I} \bigoplus A \neq \emptyset$. By Success, $\mathcal{I} \bigoplus A \subseteq [A]$, so $(\mathcal{I} \bigoplus A) \cap [A] \neq \emptyset$, and hence $(\mathcal{I} \bigoplus A) \cap [B] \neq \emptyset$. Then $\neg B \notin \mathcal{I} \bigoplus A$, and by Strong Conservativity, $(\mathcal{I} \bigoplus A) \bigoplus B \neq \emptyset$. Unfortunately, this is not quite what is needed. To plug the gap, we need a sort of converse of the desired principle – instead of the assumption that if stronger information does not crash a state, weaker information will not either, we need the assumption that if a stronger state does not crash under information, a weaker state will not either:

- **Ecnarelot**: If $(\mathcal{I} \bigoplus A) \bigoplus B \neq \emptyset$, then $\mathcal{I} \bigoplus B \neq \emptyset$.

Unfortunately, given Conjunctive update, $(\mathcal{I} \bigoplus A) \bigoplus B$ is the same as $\mathcal{I} \bigoplus (A \wedge B)$, so Ecnarelot is, in fact, equivalent to Tolerance, and no advance has been made.

However, the final principle we need for a strongly Lewis selection function – that if $[A] \subseteq [B]$ and $[A] \cap f(w, [B]) \neq \emptyset$, then $f(w, [A]) = [A] \cap f(w, [B])$ – does follow nicely from a conservativity principle. However, we need slightly more even than Strong Conservativity:

- **Very Strong Conservativity**: If $\neg A \notin \mathcal{I}$, then $\mathcal{I} \bigoplus A = \mathcal{I} \cap [A]$

This amounts to the use of Stubborn Update when Stubborn Update will not crash the state. Given Success, we already have that $\mathcal{I} \bigoplus A \subseteq [A]$, so Very Strong Conservativity then amounts to two additional assumptions:

1. We should not exclude any A possibilities from a state when updating with A.

2. We should not add any new possibilities to a state when updating it.

Again, neither of these assumptions is irresistible, but they are both not without some appeal.

Thus ends the discussion of the role of the antecedent. The story we've told here is not a novel one – the goal is not to give a new picture of the semantic role of the antecedent, but rather to use a familiar picture as a model for what a properly explanatory story might look like. The project from here – and this is the more-or-less new bit – is to turn to the consequent and try to tell a similar story there. This, we think, is something people have not much attended to. But before turning to the consequent, there's one bit of ground-setting needed.

3. Interlude

Here's a bit of potted history. Time was that the main task in picking a conditional was to select between the bargain-bin material conditional and the gold-standard strict conditional. Those who favoured the material conditional were always the rebels in this dispute, so the strict conditional emerged as the Old School conditional. The Old School resistance to the material rebels is, traditionally, one of inferential conservativism – the Old Schoolers want a conditional which licenses fewer inferences than the material conditional does (so, in particular, *not* the paradoxes of the material conditional). (The lay of the land is really more complicated than this, since the strict conditional also licenses inferences that the material conditional does not. But it is fair to say that it was the inferential (over-)strengths, rather than the inferential weaknesses, of the material conditional that were active as motivations in the dispute.) But in the middle of the twentieth century, a new dispute arose, and a new resistance to the Old School emerged. This new resistance carried the inferential conservativism of the Old School even further, wanting to reject inferences that even the Old School accepted. Paradigmatic, and historically primary, among the inferences to be rejected was Strengthening of the Antecedent, but others, such as transitivity and contraposition, soon came into question as well. Thus was born the New School conditional.

Consider a typical piece of New School departure from the Old School. Lewis (1973) gives us intuitively compelling counterexamples to contraposition, such as:

- If Boris had gone to the party, Olga would have gone.

- Therefore, if Olga had not gone to the party, Boris still would not have gone.

Setting aside the suspicious introduction of "still" into the consequent, New Schoolers hold that this inference should not be valid – as Lewis says, "suppose that Boris wanted to go, but stayed away solely in order to avoid Olga, but Olga would have gone all the more willingly if Boris had been there." The Old School cannot avoid contraposition, since once the antecedent worlds are a subset of the consequent worlds, the complement of the consequent worlds are inevitably a subset of the antecedent worlds.

The New Schoolers, of course, have a resource that the Old Schoolers lack – the selection function. The selection function can avoid the pernicious effects of the subset relation. Suppose that the worlds in which Olga goes to the party are uniformly closer than those in which she does not. Then the closest worlds in which Boris goes to the party – those picked out by the selection function for the first conditional – can be worlds in which Olga goes to the party, even though there are additional, more distant, worlds in which Boris goes to the party and Olga does not – namely, those picked out by the selection function of the second conditional (and hence counterexampling it).

But the New School triumph is hollow. Contraposition is not gone, but only buried under a thin layer of dirt. It rises to the surface with only the slightest provocation. Suppose we add the premise:

- If Olga were to either come or not come to the party, she might not come.

New School conditionals will take this premise to restore the validity of the contrapositional scheme. But surely the argument is no more valid, and the Lewisian contravening situation no less compelling, with this rather thin additional bit of information (it feels rather like a tautology, doesn't it?) added.

As with contraposition, so with strengthening of the antecedent. New Schoolers want to reject inferences such as:

- If I were to strike this match, it would light.
- Therefore, if I were to strike this match and it were wet, it would light.

Again Old Schoolers are committed to the unwanted inference – given that, in strengthening cases, the strengthened antecedent worlds are a subset of the unstrengthened antecedent worlds, if we also have that the unstrengthened antecedent worlds are a subset of the consequent worlds, it follows ineluctably that the strengthened antecedent worlds are a subset of the consequent worlds, and hence that the second conditional follows from the first. And again the New Schoolers can deploy the selection function to block the unwanted inference – here by having the *selected* strengthened antecedent worlds *not* be a subset of the *selected* unstrengthened antecedent worlds.

And again the New School triumph is hollow. So consider:

- If I were to strike this match, it would light.
- If either I were to strike this match or this match were to wet, it might be wet.
- Therefore, if I were to strike this match and it were wet, it would light.

This inference is valid according to the New School, but again the validity-restoring premise seems thin in content and insufficient to undermine the intuitive unacceptability of the inference.

Once more with transitivity. New Schoolers point out the intuitive unacceptability of many instances of transitivity, such as:

- If Bush had lost the 2000 presidential election, Gore would have won.
- If Nader had won the 2000 presidential election, Bush would have lost.

- Therefore, if Nader had won the 2000 presidential election, Gore would have won.

The validity of this inferential pattern under the Old School follows immediately from the transitivity of the subset relation. New Schoolers use the selection function to block the inference by guaranteeing that the worlds invoked by "Bush lost" *qua* consequent (namely, all of the Bush-loses worlds) are not identical to (or even a subset of) the worlds invoked by "Bush lost" *qua* antecedent (here, only the *selected* Bush-loses worlds).

And once more, the unwanted inference lurks slightly below the surface. In this case, Lewis (1973) shows how to recapture transitivity. One way is to expand the argument to:

- If Bush had lost the 2000 presidential election, Gore would have won.
- If Nader had won the 2000 presidential election, Bush would have lost.
- If Bush had lost the 2000 presidential election, Nader might have won.
- Therefore, if Nader had won the 2000 presidential election, Gore would have won.

This argument is, according to New Schoolers, valid. This particular putative validity is more familiar than the last two, and many may be ready to accept its validity, on the grounds that the first and third premises effectively contradict one another, guaranteeing the impossibility of Bush losing. Perhaps this is the right reaction, but we suspect that looking at the argument with a fresh eye shows it to be less compelling than the theoretically-inflected vision presents it as being.

The New School thus repeatedly has difficulty with the re-emergence of (modified forms of) exactly those inferential forms the rejection of which was the New School point of departure from the Old School.

> **Methodological Digression**: The situation here perhaps points to a moral for inferential consequentialism, especially of the negative variety. Suppose we are moved to acceptance of theory T because of its identification of some inferential form $\Sigma \vDash A$ as an invalid inference. But a natural language expression of $\Sigma \vDash A$ can always be regarded as enthymematic, and we know that there are *some* reconstructions of the enthymemes that result in a valid argument (minimally, those which supplement the argument with $\Sigma^* \supset A$, where Σ^* is the conjunction of some finite subset of Σ). To be moved by negative inferential consequentialism is thus to implicate oneself in an open-ended sequence

> of inferential disputes, in which one must both pass intuitive judgement on various enthymematic reconstructions, and then ensure that T agrees with the intuitive judgements. Perhaps there's no insurmountable methodological problem, because perhaps the burden thereby incurred can be met, but it does create some worries for the use of negative inferential consequentialism as a theory selection tool. In particular, the relevant enthymematic reconstructions are likely to become increasingly convoluted (since any rejection of a reconstruction leaves open the opportunity of producing enthymematic reconstructions of it), and intuitive judgements are likely to give out at some point. This may help explain the popularity of a "to the victory go the spoils" supplementary methodological principle among inferential consequentialists.

The basic structure of the New School difficulty is: New Schoolers eliminate unwanted Old School inferences via the selection function, which serves to break certain inferential connections between distinct antecedents. But so long as the selection function is reasonably structured, it will be possible to amplify the argument with additional premises which serve to reinstate the inferential connections thus broken. Each of the supplementary premises added above plays this role. As a result, the full inferential strength of the Old School conditional lies right below the surface of the New School conditional. The New School has looked in the wrong place to dampen the power of the strict conditional. The relevant inferences are, at root, a consequence of the role of the *consequent* in the evaluation of the conditional – once the antecedent has played its role, either naively (in the Old School) or via a selection function (in the New School), we look *in both cases* for a subset relation between the determined worlds and the consequent worlds. The inferential features of the Old School strict conditional thus immediately track the structural features of the subset relation – the New School conditionals also track those features, but less overtly so, since one must see past the fog created by the selection function. The New School, same as the Old School – both are driven by the subset relation.

Let's not get fooled again. If we want to make an informed decision on conditional inferential features such as contraposition, strengthening of the antecedent, or transitivity, then we need to understand the way in which those features emerge from a proper understanding of the semantic role of the consequent. It is to the search for such an understanding that we now turn.

4. The Consequent

Recall again that the Ramsey conception asks for two tasks in the evaluation of a conditional – a normatively governed transition to a new information state, and an evaluation of a further claim against that novel information state. If the normatively governed transition is to be achieved under the semantic influence of the antecedent, it is natural to look to the consequent for the performance of the second task. The Old School strict conditional then takes an appealingly simple stand on how that performance is carried out:

- **Subset Testing**: Once the antecedent has generated a particular set of worlds, the consequent's semantic function is to check if the antecedent-generated worlds are a subset of the worlds supporting the consequent. If so, the conditional is true; if not, it is false.

This picture of the semantic role of the consequent is shared by Old and New School, as is easily seen by comparing the relevant semantic clauses:

- **Old School**: $A \Rightarrow B$ is true at w iff $[A] \subseteq [B]$.
- **New School**: $A \Rightarrow B$ is true at w iff $f(w, [A]) \subseteq [B]$.

In both cases, once their divergent views on the contribution of the antecedent have been accommodated, what remains is a simple subset comparison between an antecedent-inflected body of information and the set of worlds determined by the consequent.

In the previous section, we have argued that many, and many controversial, inferential features of the conditional are ultimately rooted in this subset-based semantic contribution of the consequent. Following our general procedure, we would prefer to have a semantic picture from which we can trace the individual emergence of particular inferential features, and hence have the tools for engaging in careful conceptual consideration of which features, and which accompanying emergence conditions, are most central to the role of the conditional – rather than accepting the whole parcel in one fell swoop via the very strong subset relation. We thus begin with the following retreat to a more abstract, and less committal, location. One can think of the Ramsey conception as telling us that the consequent plays the role of imposing a test on a body of information determined by the antecedent. Thus what is needed from the consequent is a notion of *passing the test*. From a maximally abstract perspective, a notion of passing a test amounts to a partitioning of information states into the test-passing and test-failing ones. Thus we need the consequent to provide us with a set of (passing) information states. This produces the following semantic picture:

- $A \Rightarrow B$ is true at w iff $f(w, [A]) \in g([B])$ (where $g : \wp(W) \mapsto \wp\wp(W)$).

(This clause can be made Old School friendly by taking the selection function f to be the identity function on the second argument.) The strong assumption shared by Old and New School alike is thus that the function g is simply the power set operation on $[B]$ (or, to put it in terms that will later turn out to be more revealing, the map from any set to the principal ideal generated by that set). We will thus refer to both Old School and New School conditionals as *powerset conditionals*. Our central task in this section is to consider two questions:

1. What are the alternatives to a powerset conditional?

2. Why a powerset conditional, rather than some alternative?

> **A Brief Remark On Method**: The plan from here is to consider various structural constraints that can be placed on the g function, and see how these might collectively add up to the powerset conception and individually back particular inferences and be subject to particular conceptual justifications or criticisms. But the extraction of particular inferential patterns from structural features of the g function is persistently clouded by the role of the selection function (in New School conditionals). For example, one might want to say that transitivity of the conditional is backed by a corresponding "transitivity" of the g function – the requirement that if $[A] \in g([B])$ and $[B] \in g([C])$, then $[A] \in g([C])$. However, while this is straightforwardly sufficient for the Old School strict conditional, it does not suffice to produce transitivity in the New School conditional, because the switch from $[B]$ (in the consequent of the first conditional) to $f(w, [B])$ (in the antecedent of the second conditional) blocks straightforward application of the "transitivity" feature of g. We deal with this difficulty by ignoring it – we simply assume throughout the following discussion that inferential principles can be read off of structural constraints in the way possible with strict conditionals. The background assumption is that it will always be possible, using the general methods of supplementary premises serving to coordinate antecedents-post-selection-function, to retrieve some relative of the inference in question for the variably strict conditionals as well, and that persistently worrying about such retrieval will only cloud the discussion.

4.1. Van Benthem's Reduction Theorem

The principal predecessor in the literature to the sort of investigation we want to pursue is that of van Benthem (1984). Van Benthem takes a semantic theory for a conditional to require a method for mapping an antecedent (considered as a set of worlds) to a set of permissible following consequents (considered again as sets of worlds). Thus a semantics for the conditional treats the antecedent as at the same type level – but with category e replaced with the intensional category w of worlds – as generalized quantifiers. In $A \Rightarrow B$, the antecedent A generates a semantic value of type $<< w, t >, t >$, just as in "Some philosophers snore", "some philosophers" generates a semantic value of type $<< e, t >, t >$. Since A itself is of semantic type $< w, t >$, we can think of the conditional itself, thought of as the tool for mapping the antecedent to its semantic value, as of type $<< w, t >, << w, t >, t >$ – here matching the typing of "some" as $<< e, t >, << e, t >, t >$. Given this framework, van Benthem can then, following the typing analogy with generalized quantifiers, proceed to ask *which* determiners provide reasonable semantic values for the conditional.

Note that we have framed the investigation in the reverse direction from van Benthem – we take the consequent, not the antecedent, to be doing the semantic work of generating a structure against which the other part of the conditional is to be tested. (Formally, this is a matter of taste – as with generalized quantifier theory, one can take the necessary semantic feature as a relation between, or binary function on, two semantic values (restrictor and matrix, antecedent and consequent), or one can two-step the process by taking the semantic feature to be a function from one semantic value to a function on another semantic value. But we are here choosing to pursue the Ramsey-inspired hope that treating the necessary semantic feature as consequent-centered will connect it to governing intuitions on the nature of conditionality.) This means that some formal translation is needed between our way of putting things and van Benthem's. Consider, for example, our principal target: the powerset conception. Put in our way, the powerset conception takes the relevant semantic function of the conditional to be production of a principal ideal off of an input set of worlds – from the consequent we generate the set of all subsets of the consequent, and then we check whether the antecedent-delivered worlds are in that set. Put in van Benthem's way, the powerset conception takes the relevant semantic function of the conditional to be the production of a principal *filter* off of a (different) input set of worlds – from the antecedent we generate the set of all *supersets* of the antecedent, and then we check whether the consequent worlds are in that set. So in van Benthem's terminology, the powerset conception amounts to the view that the conditional is the determiner "all".

Since we are interested in conditionals less structured than powerset conditionals, it is thus of immediate interest that van Benthem lays out a short list of (what he takes to be) intuitively compelling constraints on the determiner that the conditional imposes, from which constraints it then follows that that determiner can only be selected from among:

- "All", "some or all", and "half or more"

Already, this result would rather severely constrict the space of investigation. However, things are even worse than this list of three might suggest, because a bit of thought shows that both "some or all" and "half or more" produce at best very peculiar conditionals. The determiner "some or all", for example, supports the following inference:

- **Weakening of the Antecedent**: $\Diamond A,\ A \Rightarrow B \vDash (A \vee C) \Rightarrow B$

Since "some or all" also gives $A \Rightarrow A$ as a theorem for any A, we have the following consequence:

- **Easy Connection**: $\Diamond A \vDash \top \Rightarrow A$

Anything that is possible is a conditional consequence of a tautology. We are inclined to take this as a *reductio* of "some or all"'s status as a potential conditional.

The "half or more" conditional is not so flagrantly bad as "some or all", but it still produces a number of features that look at best dubious as features of a conditional. Like all non-universal proportionality quantifiers, it fails to endorse the inference:

- $A \Rightarrow B,\ A \Rightarrow C \vDash A \Rightarrow (B \wedge C)$

And since "half or more" does not require a majority, it allows for the satisfiability of the triad:

- $\Diamond A,\ A \Rightarrow B,\ A \Rightarrow \neg B$

Thus the van Benthem result looks, for practical purposes, to reduce the playing field to a single candidate: "all", or the powerset conditional.[14]

We begin, then, with some resistance to the van Benthem reduction result. The result requires the following principles:

1. **Quantity**: Given sets of worlds X, Y, Z, and W such that:

    a) $|X \cap Y| = |Z \cap W|$
    b) $|X - Y| = |Z - W|$
    c) $|Y - X| = |W - Z|$
    d) $|\overline{X \cap Y}| = |\overline{Z \cap W}|$

[14] See Lapierre (1995) for more details on the inferential features of the "some or all" and "at least half" quantifiers.

then Y is in the generalized quantifier generated by X (henceforth, $X \Rightarrow$) iff $W \in (Z \Rightarrow)$. Hence if $[A] = X$, $[B] = Y$, $[C] = Z$, and $[D] = W$, then $A \Rightarrow B$ iff $C \Rightarrow D$.

Quantity mirrors, for conditionals, the preservation under domain isomorphism condition which is standardly taken as part of the definition of a (logical) generalized quantifier. (See Barwise and Cooper (1980), and Sher (1991) for detailed defense of the importance of Quantity.) Van Benthem recognizes that Quantity is a controversial assumption in that it presupposes that the conditional is indifferent to the *organization* of modal space, thus making it blind to (for example) a sphere-system arrangement of worlds. (Van Benthem thus runs another version of his reduction result which weakens Quantity to a condition of Quality which requires respecting further organizational features of modal space; this second version produces the Old School strict conditional and the New School Stalnaker-Lewis variably strict conditional as the only available options. More on the second version below.) But from the current perspective, this form of resistance to Quantity is misguided, since the semantic impact of the organization of modal space finds its home in the application of the selection function to the antecedent – once antecedent-worlds are selected, the required relation between them and consequent worlds is, as Quantity demands, purely quantitative.

But there is another path of resistance to Quantity. *Collective* properties are typically not preserved under isomorphism in the way that Quantity demands, and that distributive properties are. To consider an extreme case, take the sentence:

- The men surrounded the house.

If we permute the domain, keeping constant the cardinalities of the boolean combinations of *men* and *surrounds the house*, we have no guarantee of preserving truth. For the individual-level extension of *surrounds the house* is in typical cases empty, which means that the cardinality-preserving isomorphism is compatible with absolutely no surrounding occurring. The difficulty, of course, is that the relevant *surrounding* feature emerges only at the level of supra-atomic elements of the domain lattice, and these features are not respected by the relevant isomorphism. Less trivial examples, in which the predicated feature is also held at the individual level, are also available. Thus a sentence such as:

- The men are happy.

can require for its truth not just the separate happinesses of the individual men – that preserved under isomorphism – but also the way in which their individual happinesses *coordinate* to give rise to a collective group happiness.

Transferring this analogy to the case of conditionals, imposing conditions on the space of worlds, what we need is the possibility that consequents are imposing collective, rather than distributive, tests on the antecedent-inflected information states. This possibility is promoted by thinking of the antecedent precisely as delivering an information state, rather than a set of worlds (despite, of course, the formal identity of these two in the current setting) – a set of worlds cries out for testing of each world, while an information state looks for a testing *as state*, which would then be collective from the point of view of worlds.

is irrelevant to the holding of the conditional. Antecedence is the analog for conditionals of the property called by Barwise and Cooper (1980) *living on* – a quantifier lives on its restrictor if objects which do not satisfy the restrictor are irrelevant to the applicability of the quantifier to a matrix.

Van Benthem's defense of Antecedence is rather brief – he says only that "a conditional statement invites us to take a mental trip to the land of the antecedent" (311). Note first that this line of defense requires the assumption that the selection function obey the Success-derived constraint $f(w, [A]) \subseteq [A]$ – without this constraint, the Antecedence requirement can allow the relevance of non-antecedent-satisfying worlds in the evaluation of the consequent.

Again a collective approach to the semantics can call Antecedence into question. How things stand with a group can depend not only on how each of those things stand, and also not only on the (Quantity-conflicting) relations among those things, but also on how other things, brought into consideration by the move to the collective perspective, stand. Thus whether a neighborhood is beautiful may depend not only on whether each house in it is beautiful, but on whether additional houses along the neighborhood boundary (but not in the neighborhood) are beautiful. (Note that this can be true even if these additional houses are not relevant to the beauty of any individual neighborhood house, even those on the boundary). This anti-Antecedence effect *can* be achieved via the antecedent, through the selection function, but one can also follow a particular line of conception for the selection function and still think there is need for an additional anti-Antecedence effect.

2. **Activity**: Van Benthem argues that "a logical constant should do some work, showing some variety of behaviour within its proper field of action" (312-313), and on the basis of this consideration requires conditionals to satisfy the requirement of Activity. Activity demands that for any non-

empty antecedent, there be both a consequent that produces a true conditional in combination with that antecedent, and also a consequent that produces a false conditional in combination with that antecedent.[15]

Van Benthem's stated justification for Activity is not terribly compelling. The principle would, for example, fail for both (truth-functional) conjunction and disjunction. One might start with the thought that a unary logical connective should show activity, and then generalize to the n-ary case, but the relevant generalization seems to be that there be some n-tuple which the connective maps to the true and some n-tuple that it maps to the false. *That* constraint is easily met without meeting Activity. Also, some justification for picking empty antecedents (and full consequents, if we generalize as in the previous footnote) as an exception to the Activity requirement seems wanted.

We think the case for Activity can in fact be made more convincingly from the perspective of the Ramsey test. From this perspective, Activity amounts to the assumption that any information state will pass some tests and fail others. It is here tempting to drop van Benthem's restriction to non-empty antecedents, and assume that the relevant updating procedure will return a coherent information state, even if the input information is incoherent. One might then take this constraint to be more or less constitutive of the notion of information – as something which splits possibilities, and which thus rules some things in and other things out.

3. **Confirmation**: Van Benthem takes Confirmation to mark out the "most distinctive feature of conditionality" (311). The core of Confirmation tracks the way that conditional truth is updated under reception of supporting and disconfirming cases. We begin with the thought that a *positive instance* is a possibility which makes true the consequent, and a *negative instance* is one which makes true antecedent but not consequent. (It is not clear why van Benthem omits from the notion of positive instance cases which falsify the antecedent.) The basic governing principle is then that neither adding positive nor removing negative instances should remove truth. Van Benthem breaks this down into three cases:

   a) Adding Pure Positive Cases: Suppose $A \Rightarrow B$ is true, and we add the ideal positive case: a possibility which supports both antecedent and consequent. Then truth should be preserved. Thus from $A \Rightarrow B$, we can infer $(A \vee C) \Rightarrow (B \vee C)$.

[15] The same conceptual motivation would seem to back another constraint that van Benthem does not impose, requiring that for each non-tautologous consequent, there be some antecedent with which it forms a true conditional and some antecedent with which it forms a false conditional.

b) Adding Incidental Positive Cases: Suppose $A \Rightarrow B$ is true, and we add a merely incidentally positive case: a possibility which supports consequent but not antecedent. Then truth should be preserved. Thus from $A \Rightarrow B$, we can infer $A \Rightarrow (B \vee C)$.

c) Removing Negative Cases: Suppose $A \Rightarrow (B \wedge C)$ is true, and we remove a negative case (let the negative case be a C possibility). Then truth should be preserved. Thus from $A \Rightarrow (B \wedge C)$, we can infer $(A \wedge C) \Rightarrow B$.

The distributive, rather than collective, conception of the conditional is on display again in this picture of Confirmation in the formal realization of the intuitive notions of positive and negative instances. Modulo that concern, these principles do seem to spring deeply from the heart of the Ramsey conception. Van Benthem then adds one final dimension to Confirmation: he requires that "optimal evidence" suffice for the truth of a conditional, and hence that $A \Rightarrow B$ be true whenever $[A] \subseteq [B]$. This final constraint seems to us to flow less naturally from the underlying thought of Confirmation – it represents, as it were, a boundary condition on the normative processes invoked by the Ramsey conception, rather than an integral part of the functioning of that process.

4. **Uniformity**: Van Benthem formulates one final – and, as we will see below, very powerful – constraint on the logic of the conditional. The driving thought is that we ought to be able to "build one's way" from the truth value of a conditional on a small scenario to the truth value of a conditional on a large scenario, via discrete steps. More precisely, suppose one knows the truth value of $A \Rightarrow B$, and then one proceeds to expand the scenario in one of three ways:

   a) Adding a positive instance

   b) Adding a negative instance

   c) Adding both a positive and a negative instance

   Then the principle of Uniformity says that the way in which the result of expanding in the third way depends on the starting truth value and the way of expanding (severally) in the first and second ways must be fixed for all conditionals.

Van Benthem provides the following useful graphical conception of Uniformity. Arrange the original truth value and the three expansion truth values in a diamond formation such as:

$$\begin{matrix} & 1 & \\ 1 & & 0 \\ & 1 & \end{matrix}$$

This diamond represents the claim that a true conditional is such that (a) when a positive instance is added, it remains true, and (b) when a negative instance is added, it becomes false, and (c) when both positive and negative instances are added, it remains true. Uniformity then rules out the diamond:

$$\begin{matrix} & 1 & \\ 1 & & 0 \\ & 0 & \end{matrix}$$

in which the same first three cases get mapped to a different fourth outcome.

Uniformity thus drops the *a priori* space of sixteen possible construction techniques down to merely eight. Some of these are in turn ruled out by the principles of Confirmation, leaving only four permissible methods (although which four will depend on which of the incompatible pairs of diamonds are selected).

Even if one accepts the broad general thought behind Uniformity, the particular realization seems *ad hoc* in its details. Most pressing among the details is the particular choice of supervenience base. The base seems simultaneously too broad and too narrow:

- Too broad, because it is not clear why the *original* truth value of the conditional is relevant to calculating the truth value resulting from simultaneous addition of positive and negative instances, given that we already have in hand the several results of adding a positive instance and of adding a negative instance.
- Too narrow, because it is not clear why we want only to consider transitions from single-case additions to dual-case additions. Why not also, for example, think that novel patterns emerge when we consider simultaneously adding two positive and a negative case, which patterns cannot be reduced to patterns emerging at the dual-addition level?

The second of these worries in particular shows that Uniformity serves to rule out a broad range of *proportionality* quantifiers as candidates for conditional semantics. Suppose we wanted that a conditional be true so long as 90% of the antecedent-selected instances were positive instances. This pattern cannot be resolved into a Uniformity-supporting sequence of individual and dual additions – adding a positive and a negative case simultaneously is truth-supporting *provided that we had a sufficient excess of positive cases initially*. But this fact will not be revealed in the supervenience base for Uniformity. A broader version of Uniformity, in which the supervenience base was extended to the results of adding mixtures of nine cases and the supervening cases were the results of adding mixtures of ten cases, *would* permit this proportionality quantifier (but not others) as a conditional.

With these five principles is hand, we can now state and prove van Benthem's reduction result:

Theorem (van Benthem (1984))**:** The only conditionals-as-quantifiers on finite domains which satisfy Quantity, Antecedence, Activity, Confirmation, and Uniformity are *all*, *some or all*, and *at least half*.

Proof: The role of Quantity and Antecedent is restricted to allowing a convenient modelling of the situation. Consider a conditional $A \Rightarrow B$. Antecedence tells us that only the sets $[A] - [B]$ and $[A] \cap [B]$ matter to the truth value of the conditional. Quantity then tells us that only the cardinalities of these sets matter. So the full range of cases can be modelled – given that we are concerned only with finite domains – as ordered pairs of natural numbers representing the cardinalities of $[A] - [B]$ and $[A] \cap [B]$. These ordered pairs can then be arranged in a tree structure:

$$
\begin{array}{ccccccc}
 & & & 0,0 & & & \\
 & & 1,0 & & 0,1 & & \\
 & 2,0 & & 1,1 & & 0,2 & \\
3,0 & & 2,1 & & 1,2 & & 0,3
\end{array}
$$

A choice of quantifier then amounts to an assignment of truth value to each node in this tree.

We now start tracing out constraints on the available trees:

1. The top 0,0 node must be marked 1, due to the fourth part of Confirmation. Note that it is this fact and this fact only that prohibits *some* from being a quantifier satisfying van Benthem's five constraints.

2. The second row (containing 1,0 and 0,1) must be marked 1 0, due to Activity (which tells us that there must be a true entry and a false entry) combined with the fourth part of confirmation (which again tells us that when all cases are positive instances, the conditional is true). Thus the tree must begin:

   1

   1 0

3. By Confirmation, adding a positive instance always preserves truth. Thus there can be no diamonds in the tree of the form:

   1

   0 $*$

   $*$

   (where $*$ marks an arbitrary truth value). Similarly, adding a counterexample must preserve falsity (this by applying the third part of Confirmation in reverse), so we can have no diamonds of the form:

   0

   $*$ 1

   $*$

   Now suppose we start with a truth, and that adding a counterexample preserves truth. Then adding both must also preserve truth (since this amounts to a truth-preserving addition of a positive instance to the truth created by adding a negative instance). This means that truth is down-and-left monotonic in the tree. Similarly, if we start with a falsity and that falsity has two falses below it in the diamond, the result of adding both positive and negative instances must be false. So falsity is down-and-right monotonic in the tree.

   Combining these results tells us that the tree in general has the form of a zig-zag boundary layer between the true and the false, with the true to the left of the boundary and the false to the right of the boundary.

4. Now Uniformity sharply constricts the shape of the zig-zag boundary line. The shape of the boundary line depends on whether the node directly under the 1-to-0 meeting point of one line itself becomes 1 or 0. But, given Uniformity, this means there are only two cases to consider. At any given point, to choose the next zig or zag is either to be in the situation:

    1

    1 0

    *

    or in the situation:

    0

    1 0

    *

    Uniformity tells us that only one choice can be made, for the entire tree, in each of these two situations.

    Suppose we choose in the first situation to complete the diamond with 0. Then the second situation will never arise – from here on, the boundary point always lies below a 1, and the rest of the tree is determined. This produces a tree corresponding to *all*.

    Suppose we choose instead in the first situation to complete the diamond with 1. The very first boundary point decision is of the first situation type, and completing it with 1 then forces the second boundary point decision to be of the second situation type. If we also complete this with 1, then all subsequent boundary decisions are of the second type and are completed 1. The result is a tree corresponding to *some or all*.

    Suppose instead that we complete the second situation type with 0. Then the location of the boundary henceforth bounces back and forth between the first situation type and the second situation type, and the boundary line thus zigs and zags repeatedly. The result is a tree corresponding to *at least half*.

Since this exhausts the decisions left open by Uniformity, we have thus characterized all available trees, and all available quantifiers to use in giving the semantics of conditionals on finite domains.

We have treated van Benthem's proof in some detail on the general methodological grounds that impossibility, triviality, and reduction results are valuable philosophical commodities and ought to be carefully appreciated – even where the premises are subject to serious doubt, such results serve to shape investigation and illuminate the logical landscape. As our glosses on Quantity, Antecedence, Activity, Confirmation, and Uniformity have indicated, we think there are interesting directions of conceptual development which evade the pressure toward van Benthem's various constraints, and we thus do not take the reduction result to be successful in foreclosing the logical space from which quantifiers can be selected as conditionals (in fact, we do not even think it necessary that the conditional semantic operation take on a quantifier-like form – more on this below). But the result does help show *how* one must conceive the conditional in order to make room for anything other than the van Benthem Three (and, if our earlier observations on the unacceptability of *some or all* and *at least half* are correct, really the van Benthem One).

Some more specific observations:

1. The central role of Uniformity in the result should be emphasized (and, indeed, *is* emphasized by van Benthem). Without Uniformity, we get the substantially weaker result that any zig-zag boundary path produces an acceptable quantifier for the conditional, yielding 2^{ω} rather than three choices. Since Uniformity strikes us as by a substantial margin the least-well-motivated of van Benthem's constraints, its central role is particularly troubling.

2. It is worth considering the way in which some quantifiers of interest are foreclosed by van Benthem's result. So:

    a) One might be tempted by a strong proportionality quantifier, having the thought that it suffices for $A \Rightarrow B$ that the vast majority of A worlds are B worlds. But strong proportionality quantifiers are always defeated by Uniformity since, as observed above, the acceptability of adding simultaneously a positive and a negative instance will depend on the extent by which the proportionality threshold is currently exceeded, and thus cannot be extracted from Uniformity's supervenience base. Since this fact is directly a consequence of what looks like an *ad hoc* decision about that supervenience base, the resulting elimination of a tempting conditional semantics seems less than persuasive.

b) One might be tempted by an *all but one* quantifier, or some variant thereof, again on an error-tolerant conception of the conditional. But *all but one* fails of Activity, since it produces truth in both cases when the domain size is 1. Were we to try to evade the difficulty in the small domain size by transitioning at some domain size from *all* to *all but one*, we would fall afoul of Uniformity.

c) The quantifier *some*, of course, does not seem a tempting choice for the conditional (unless one is interested in a "might"-conditional), but it is a curiosity of the van Benthem result that it permits the odd *some or all* but not the much more standard *some*. The reason lies entirely in the fourth clause of Confirmation, which we have already observed to be less deeply motivated than the rest of Confirmation. This clause requires every tree to be rooted at 1, whereas any existentially committing quantifier such as *some* roots the tree at 0.

d) Van Benthem's result requires that the domain be finite. This assumption can seem unexceptionable – we can, after all, directly generate a finite space of worlds by considering all assignments to the sentence letters used in the conditional(s) under consideration. (Issues become more complicated where quantification is involved, but we restrict ourselves to the easiest cases for now.) But again the distributive, rather than collective, impulse seems to be required to justify this move. In particular, if we deny Antecedence, then we open the possibility that worlds other than antecedent-supporting worlds are relevant to evaluation of the conditional, and this in turn raises the question of which features of those worlds determine the nature of their relevance. If the antecedent itself is not doing the determination work, then it may be that that work is not strictly syntactically constrained by the conditional, and hence may require more articulation of modal space than the conditional's syntactic material provides. (There is more to be said here, as van Benthem's adaptation of his reductive result to the infinite domain case needs careful consideration.)

4.2. MUDdy Conditionals

While we are skeptical of a number of van Benthem's constraining principles, and hence skeptical of the specific final result, we are sympathetic with the general exercise of laying down well-motivated structural constraints and observing the way in which these constraints shape the range of available conditional options. Our interest will be more in surveying the options than in selecting from or limiting those options, however.

We lay out three types of constraints here. The first collection of constraints serve to shape the algebraic structure generated by a consequent. In their strongest form, these constraints will generate a boolean algebra. The shape of the algebraic structure will be seen primarily to control (or constrain) inferences between conditionals with the same consequent and distinct antecedents. The second collection of constraints serve to relate the algebraic structures generated by distinct consequents; these constraints then control/constrain inferences between conditionals with the same antecedent and distinct consequents. The final collection of constraints then locates this entire family of inter-related algebraic structures, which absent these constraints can float freely within the power set of the power set of the space of worlds.

> **A Preliminary Question: Are Conditionals Quantifiers?** The question of whether conditionals are quantifiers, as van Benthem takes them to be, amounts (again following the lead of Barwise and Cooper (1980) to the question of whether they satisfy Quantity and Antecedence from above. Since Quantity is, in essence, a question of whether the algebraic structure generated by a consequent has its cardinality features fixed by the cardinality of the consequent, the Quantity question can be affirmatively answered by (a) pinning down the structural nature of the relevant algebra tightly, and then (b) pinning down the location of that algebra relative to the consequent itself. Thus the strongest versions of the conditional that we contemplate below (such as the powerset conception) will automatically answer the Quantity question affirmatively. Independent of such automatic answering, we can see no particular reason to agree to Quantity, nor much independent theoretical interest in it. Antecedence we take up in more detail in the second section below.

4.2.1. *Selecting An Algebra*

We begin with the constraint that the consequent generate a condition which is *downward closed*:

- **D**: If $X \subseteq Y$ and $Y \in g(Z)$, then $X \in g(Z)$

The D constraint corresponds to the thought that the conditional provides an operation (quantificational, if appropriate) which is *antipersistent* in the sense of Barwise and Cooper (1980).

Downward closure backs the inferential feature Strengthening of the Antecedent:

- **SA**: If $A \rightarrow B$ and $C \vDash A$, then $C \Rightarrow B$.

SA is, of course, famously objectionable from the New School perspective. But one must be cautious: New School techniques can accept D and still avoid the crudest form of SA, because of the role of the selection function. But, following on the moral of the Interlude above, some form of SA will remain deeply written into the New School conditional so long as it runs on an algebra obeying D.

To endorse D is to commit to the consequent generating a maximal partial ordering under the subset relation. If one thinks of the consequent as imposing a test on the antecedent information state, then D represents a natural thought, since it encapsulates the idea that an enriching of the tested information state leads to an increase in the number of passed tests.

From D we automatically derive that the generated algebraic structure is closed under intersections, and thus supports:

- **I**: If $X \in g(Z)$ and $Y \in g(Z)$, then $X \cap Y \in g(Z)$.

This follows from the fact that $X \cap Y \subseteq g(Z)$ (and shows that the premises of I are unnecessarily strong in the presence of D. But this way they correspond to a natural algebraic constraint). I backs the following inference pattern:

- **Antecedent Conjunction**: If $A \Rightarrow B$ and $C \Rightarrow B$, then $(A \wedge C) \Rightarrow B$.

Again this principle is not New School acceptable in its pure form, but a selection-function-conditioned version of it is.

Corresponding to closure under intersections is closure under unions:

- **U**: If $X \in g(Z)$ and $Y \in g(Z)$, then $X \cup Y \in g(Z)$.

Closure under unions does not follow from D. If both I and U are imposed, then the consequent-generated structure is a lattice, and if (stronger) both D and U are imposed, the structure is an ideal. The latter structure is suggestive, since ideals play a number of interesting logical and mathematical roles – in particular, they dual with filters, which serve as an algebraic generalization on the notion of a proposition. The D principle has a naturalness to it under the Ramsey conception, since it amounts to the assumption that if the consequent imposes a test passed by two information states, it also imposes a test passed by the meet of those two information states. This in turn corresponds to the legitimacy of proof by cases as a mode of reasoning in ascertaining whether the consequent test is passed.

The principle U backs the following inference pattern:

- If $A \Rightarrow C$ and $B \Rightarrow C$, then $A \vee B \Rightarrow C$

This principle is, in fact, endorsed by New School conditionals.

Given both D and U, we obtain a (distributive) lattice sitting inside the power set of the power set of the space of worlds (and hence containing a minimum and, assuming finiteness, a maximum) and forming an ideal within that larger boolean algebra. But we still have less algebraic structure than the full powerset conception gives us, since that conception has each consequent generate a boolean algebra. To bring us up to that level of structure, we need one more constraint:

- **N**: If $X \in g(Y)$, then $\overline{X} \in g(Y)$

There is at this point an essential underspecification in the principle N – we want that if a set is in the structure generated by a consequent, then *some sort* of complement of that set is also in the structure. Thus we need a background set against which to calculate the complement. We might, but need not, use the full space of worlds, or the set Y itself. Whichever set we pick will automatically become the maximum in the lattice (on the assumption that we endorse principle U).

N is, we think, intuitively a rather curious principle, since it claims (more or less) that if an information state passes a test, the (perhaps relative) negation of that information state also passes that test. It thus supports an inferential principle of the form:

- If $A \Rightarrow B$, then $(C \wedge \neg A) \Rightarrow B$, for appropriate (perhaps tautological) choice of C.

N thus does not seem to us a principle that would be naturally selected by any conception of the conditional – but it is built in to the powerset conception.

We thus have a number of structural options for the algebra built off the consequent:

1. With just D, we obtain a maximal partial ordering.
2. With I and U, we obtain a distributive lattice.
3. With D and U, we obtain an ideal within a distributive lattice.
4. With I, U, and N, we obtain a boolean algebra.
5. With D, U, and N, we obtain a maximal boolean algebra.

The powerset conception, of course, moves us immediately to the strongest (D,U,N) of these structural positions.

4.2.2. *Relating Algebras*
By selection among the above principles, a theory of the conditional can fix what kind of algebraic structure each consequent generates. From this point, we can raise additional questions about how the structures generates by distinct consequents relate. One obvious such principle – the one which we take to be most compelling is a principle of *monotonicity*:

- **M**: If $X \subseteq Y$, then $g(X) \subseteq g(Y)$

Absent M or other constraining principles, one could adopt all of D, U, and N above, and have each consequent generate a maximal boolean algebra, but still have no control over the relation between (e.g.) the boolean algebra generated by A and that generated by $A \wedge B$. With M, we fix that $A \wedge B$ generate a smaller structure than A.

From the Ramsey test point of view, M amounts to the assumption that a logically stronger consequent imposes a test which is met by fewer information states than a logically weaker consequent. M gives rise to the inferential feature of *weakening of the consequent*:

- **W**: If $A \Rightarrow B$ and $B \vDash C$, then $A \Rightarrow C$.

Given the minimal inferential picture of ⇒I and ⇒E, W can appear an inevitable principle (although we return below to a route for resisting it): it amounts to the thought that we can proceed within a conditional subproof via standard logical reasoning.

The principle M amounts to the claim that the determiner giving the logic of the conditional is a monotone increasing quantifier (in the sense of Barwise and Cooper (1980)). M is also equivalent to that portion of van Benthem's Confirmation which requires that adding positive cases preserve truth. One could, of course, also consider the use of monotone *decreasing* determiners in conditional semantics. Monotone decreasing quantifiers would give the result:

- If $X \subseteq Y$, then $g(Y) \subseteq g(X)$

In general, monotone decreasing quantifiers are inimicable to the core idea of the conditional, since they ask that logically *stronger* tests be *easier* to pass. However, one limited version of a principle of this general form has some appeal. This is a version which puts in place the *living on* constraint:

- **LO**: If $X \in g(Y)$, then $X \in g(X \cap Y)$

LO, given the earlier tacit commitment to Quantity, would make the conditional a species of quantifier. LO is also equivalent to van Benthem's Antecedence, and is, of course, subject to the same worries as that principle.

From a Ramsey test point of view, LO amounts to the assumption that an information state can *self-deploy* in attempting to pass a test – so that if the state passes the Y test, it can also pass the test of Y conjoined with itself. (Again, we are here suspending from questions regarding the selection function, so this self-deployment principle is not threatened by a failure of Success, since it is the selected information state, not the antecedent proper, which must be self-deployed.) One might already worry about a principle of self-deployment on the sort of resource-sensitive ground that motivate linear logic (see Girard (1987)); we will return below to further, and related, considerations against self-deployment.

LO produces an inferential principle of *antecedent absorption*:

- **A**: If $A \rightarrow B$, then $A \rightarrow (A \wedge B)$

M and LO are both consequences of the powerset conception – M because the power set of a subset is a subset of the powerset of the superset; LO because if X is in the power set of Y, then X is a subset of Y, and $X \cap Y$ is just X, and X is in its own power set. Another tempting principle relates the algebraic structure generated by a consequent to that generated by the negation of that consequent:

- **CCP**: If $X \in g(Y)$ and $X \in g(W - Y)$, then $X = \emptyset$

CCP then produces the "exclusion" inferential pattern noted earlier *not* to be produced by the *at least half* determiner:

- **E**: If $A \Rightarrow B$ and $A \Rightarrow \neg B$, then $\Box\neg A$

E is straightforwardly produced by the powerset conception, and also, we think, flows naturally from the Ramsey conception (here it represents more or less the thought that no non-trivial information state supports contradictory conclusions – a natural, if not irresistible, thought).

A final coordinating principle produces a (higher-order) lattice structure among the algebraic (and perhaps lattice) structures already associated with each consequent. This principle relates the structures generated by two consequents to that generated by their conjunction:

- **CI**: $g(X) \cap g(Y) \subseteq g(X \cap Y)$

In the presence of M, this can be strengthened to an identity, and thus we have a join operation on the collection of consequent algebraic structures. CI backs the inferential pattern of *consequent conjunction*:

- **CC**: If $A \Rightarrow B$ and $A \rightarrow C$, then $A \Rightarrow (B \wedge C)$

M together with CI produce a lattice of consequent structures, and E then allows that to become a complemented lattice. LO is difficult to fit into this general higher-order structuring picture, since it coordinates only one specific element between distinct consequent structures, rather than a stronger subset relation. Again, the powerset conception delivers immediately to us the strongest possible position, delivering M and LO, as noted above, as well as E (since disjoint sets share only the empty set in their power sets) and CI (any set which is a subset of X and a subset of Y is also a subset of their intersection).

4.2.3. *Locating Algebras*

On the picture we are investigating, each consequent determines a structure dictating what information states pass that test. Given this picture, one thing we can ask is what each such structure looks like – here the available answers have ranged from a partial ordering to a maximal boolean algebra. Another question we can ask is what the full assortment of such structures look like. Again a range of answers are available, from a (higher-order) partial ordering (with M) to a (higher-order) complemented lattice (with M, CI, and E). But all of these investigations still leave unanswered the question of *which particular* structures are assigned.

Even if we opt for the maximally structured positions, this final locative question is wholly unanswered. There are many boolean algebras on the space of worlds, so to say that each consequent selects a boolean algebra is not yet to say *which* boolean algebra. There are even many boolean algebras of boolean algebras, so to say that all of those (lower-order) boolean algebras collect into a (higher-order) boolean algebra is still not to locate any of them. Thus we consider a final class of principles, which we think of broadly as *boundary conditions* on the g function, to answer this floating locative question.

The obvious approach to answering the locative question is to answer it from two ends – by bounding the relevant structure from below and from above. Thus we start with a principle of *anchoring*:

- **AN**: $X \in g(X)$

Anchoring then yields the principle of idempotence, or $A \Rightarrow A$. We note in passing that the features relevant to the holding of idempotence on both the antecedent end – Success – and the consequent end – anchoring – are boundary conditions, rather than deep structural conditions. This fact helps drive our (admittedly inchoate) suspicion that idempotence does not represent a conceptually central aspect of conditionality.

Anchoring bounds the algebraic structure generated by a consequent from below, requiring that that structure contain *at a minimum* the set of consequent worlds. From anchoring together with CI, we can derive LO. Suppose $X \in g(Y)$. By AN, $X \in g(X)$, and by CI these collectively entail $X \in g(X \cap Y)$, as LO requires.

Anchoring's theoretical motivation, on the Ramsey conception, lies in the thought that if there's one information state that is guaranteed to pass the test imposed by a consequent, it's the information state represented by that very consequent. This is certainly a thought with considerable appeal; our resistance to it will have to await the more careful discussion, in the next and final section, of what it is for a consequent to impose a test.

Anchoring and LO are intertwined in ways that seem to deserve further consideration. We have seen already that AN together with CI yields LO; it is also the case that LO can with plausible supplementary principles produce AN. Consider another, rather minimal, boundary condition of *seriality*:

- **S**: Given any set X, there is a Y such that $X \in g(Y)$

Seriality requires that every information state pass some test; it can thus be thought of as encapsulating half of van Benthem's Activity requirement (without the restriction to non-trivial information states). LO, given boundary condition S together with M, entails anchoring. Take a set X. By S, there is Y such that $X \in g(Y)$. Then by LO, $X \in g(X \cap Y)$. Since $X \cap Y \subseteq X$, by M we have $g(X \cap Y) \subseteq g(X)$, and hence $X \in g(X)$, as AN demands.

Anchoring imposes a lower bound on the structure generated by X, so all that remains is an upper bound. Here we want a constraint of *limitation*:

- **L**: $\bigcup X \subseteq X$

Limitation thus requires that no state that passes the consequent test contain any worlds that are not in the consequent proposition. Limitation can also be motivated on Ramseyian grounds, on the thought that non-consequent worlds in the tested information state will suffice to fail the test. The line of thought here is, we think, much less persuasive than that backing anchoring, since it requires a notion of error-intolerant testing. If a test can be passed so long as exceptions are not "too numerous", then L will seem unmotivated.

Given L, any conditional whose antecedent-selection encompasses more worlds than its consequent will come out false. It is tempting to read off of this fact an inference principle of the form:

- $\vDash \neg((A \vee B) \Rightarrow A)$

but this would, of course be incorrect (fortunately, since such conditionals can appear true). A corresponding principle is available only when there is a cross-model *guarantee* that the antecedent will outrun the consequent. Such a guarantee is available when the two are contradictories and the antecedent is not trivial, so we obtain:

- $\Diamond A \vDash \neg(A \vDash \neg A)$

If we accept Lewis (1973)'s duality between *would* and *might* conditionals, and read $\Rightarrow$ as a *would* conditional, this becomes:

- $\Diamond A \vDash A \Diamond\!\rightarrow A$

Another form of idempotence thus follows from L – the two boundary conditions effectively bookend the conditional with two commitments to idempotence.

As usual, the powerset conception delivers all the constraints at once. If $g(X)$ is just the power set of X, then g meets anchoring, seriality, and limitation. This fact lies at the heart of our resistance to the uncritical assumption that the consequent imposes its semantic function via the powerset conception – *perhaps* this is indeed the right final picture (although we are skeptical), but it moves too fast, and imposes every single structural and inferential feature in one fell swoop, without pausing to consider the individual sources and motivations of distinct such features. We hope that at a minimum, the articulation of the consequent function into distinct stages as we have sketched here will illuminate why each of different inferential features might be possessed by the conditional, and aid in thinking carefully about the relative acceptability of different such features.

The strength of the boundary conditions should be noted. From limitation alone, it follows that $g(X) \subseteq \wp(X)$, and from anchoring together with downward closure, it follows that $\wp(X) \subseteq g(X)$. Thus if one adopts both boundary constraints AN and L together with D, one is straightway committed to the full powerset conception, and hence all of U, I, N, M, LO, E, and CI. While the consequent structure floats free, one thus has considerable choice about the degree of structure it imposes, but once it is pinned in place by the boundary conditions, even quite minimal structural constraints push it all the way up to the full powerset conception.

Our stance in this section has been primarily descriptive, but we indulge in a brief prescriptive moment to close in noting that we find the less structured point combining M, U, and D to be a natural point for the consequent structure to rest.

4.3. Neighborhood Conditionals; Conditionals and Reasoning

Our concern throughout has been to trace the semantic development of the conditional from an underlying conceptual picture of what conditionality is. In the first part of this paper, we have criticized a number of pictures of the semantic role of the antecedent for insufficient attention to such a developmental archaeology, but one might reasonably feel that our own examination of the semantic role of the consequent falls short in this respect. One of our points is that the very possibility of a substantive semantic story regarding the function of the consequent has been neglected in the literature, and hence that the state of play is relatively primitive in this area. Given this, we are reasonably content just to point out the option for saying more, and to highlight some of the routes available should a compelling conceptual story be told. But we will close by considering briefly two conceptual pictures that we think at least help guide some selection of a good story of the semantic role of the consequent:

1. **Neighborhood Semantics**: We have alluded repeatedly to the possibility that the semantic function of the consequent could be a collective, rather than a distributive, predication of the information state generated by the antecedent. To pursue this possibility is to employ for the conditional a *neighborhood* semantics of the sort employed by Scott (1970) and Segerberg (1971). Some of the specific points of resistance to standard pictures (including the powerset conception) that drop out of a neighborhood perspective have been traced out above, but very generally we note that:

    a) On the antecedent side, if our interest is in picking out a modal region/information state which *as a neighborhood* is an antecedent-selected one, we may be less deeply concerned with the constraint of Success. So long as our perspective is distributive/pointwise, it is natural to think that – given an antecedent A – an A-selected information state must be *thoroughly* A, and thus consisting only of A worlds. But once the perspective becomes collective, we can note that a *region* can be an A region for reasons other than the thoroughgoing A-ness of points within that region. Indeed, we may think that univocal pointwise A status is neither necessary nor sufficient for the antecedent-selected region to be collectively A. To analogize: that every house in a neighborhood is beautiful is not sufficient for the neighborhood to be beautiful, since the individual beauties may not cohere on the collective level, and that every house in a neighborhood is beautiful is not necessary for the neighborhood to be beautiful, since beauty, like other collective features, can be emergent.

b) On the consequent side, if we think of the consequent as collectively predicating the region of worlds selected by the antecedent, it becomes natural to think that the predication may involve features beyond the world-by-world features given pointwise in the predicated region – the modes of organization of the region, and the way in which the region is embedded in its (modal) surroundings may also be relevant.

The danger with the neighborhood perspective, obviously, is that the collectivizing move proves such a powerful tool for undermining the justification of inferential patterns that we are left with no interesting logical structure for the conditional. Guarding against that danger requires further investigation of the nature of collectivity best suited for modelling conditionality, and a bringing into engagement of conceptual issues lying behind both conditionals and plurals. All of this remains a project for another time.

2. **Inferential Guarding**: We have articulated two components of the Ramsey conception of the conditional: the shaping of a body of information under the normative influence of the antecedent, and the testing of that body of information via the consequent. Given this articulation, it is natural to think of the role of the consequent as being the specification of what counts as the passing of the test. A further thought is then that what is needed to complete the semantics of the consequent is a selection of *modes of deduction*, so that one can determine whether those modes admit a transition from the antecedent to the consequent. The powerset conception then corresponds to the selection of classical logical consequence as the sole means of deduction, and resistance to the powerset can be framed as a preference for other modes.

This departure from the purely classical can happen in one of two ways. In the simpler form of departure, a theory of the conditional makes a *for-all-times* selection of a particular non-classical mode of inference, and takes the truth of the conditional to depend on the derivability of the consequent from the antecedent-shaped information under this privileged mode. Thus, for example, Gillies (2006) can be taken as recommending the use of Veltman (1996)'s *update-to-test* consequent relation as the for-all-times inferential mode in conditionals. A preference for modes of inference stricter than the classical (such as various forms of relevance logic) will cause the consequent to generate a structure smaller than the powerset conception demands (in particular, D will typically fail for such logics), while a preference for modes of inference looser than the classical (such as the deployment of abductive inference) will cause the consequent to generate a structure larger than the powerset conception demands (in particular, L will typically fail in the presence of abduction).

The simpler form of departure leaves no room for a special role of the *particular* content of the conditional to shape the available mode of deduction – on the simple form, all the particular content does is specify the boundary points of the deduction which must be carried out using an independently-specified mode of deduction. A less simple form of departure from the classical allows that the contents of the conditional themselves can serve to indicate how the inference is to be carried out. The less simple view can be approached either via example or via principle. By example, consider sentences such as:

- If we set aside worries regarding empty terms, the main object to direct reference is its inability to explain the role of belief contents in producing action.[16]
- If we restrict ourselves to intuitionistically valid modes of reasoning, then it can be that $\neg\neg p$ without it being that p.
- If B and $A \supset B$ and we allow affirming the consequent, then A.

By principle, consider the general structure of a conditional subproof procedure in a natural deduction system. In a standard formulation of such subproofs for a Stalnaker-Lewis conditional, when attempting to prove $A \Rightarrow B$ the subproof begins by taking on board the new piece of information A. Within the subproof, one can then proceed via standard deductive inference, and one can also import consequents of external conditionals whose antecedents are, or are conditionally equivalent to, A. We can thus separate two different roles of the antecedent in shaping the subproof:

a) **Informational Role**: The antecedent provides a piece of information that can be utilized in attempting to derive the consequent.
b) **Inferential Guarding Role**: The antecedent controls what sort of information external to the subproof can be imported into the subproof.

In the proof procedure for the Stalnaker-Lewis semantics, these two roles are carried out equally by the antecedent.[17] In a material or a strict conditional, the roles come apart in one way, with the antecedent playing the

[16] Thanks to George Bronnikov for bringing examples of this sort to our attention.
[17] It is, admittedly, a fact that sits ill at ease with the way we have structured the discussion that these informational and inferential guard roles are played by the antecedent, rather than the consequent (where we have preferred to place the performance of the second half of the Ramsey test).

informational but not the inferential guard role (with the material conditional, there is no inferential guard, and with the strict conditional, the inferential guard is uniformly a barrier to non-modal material, independent of the conditional at hand). But we could also envision conditionals in which the two roles came apart in the other direction – conditionals in which the antecedent placed local constraints on what inferential procedures were available without supplying itself as a piece of information to be deployed in the subproof. If we think of the antecedent in part as the gatekeeper of the conditional subproof, the acceptability of a principle of idempotence becomes the question of whether the bouncer always admits himself to the club he keeps.

References

Adams, E.: 1970, 'Subjunctive and Indicative Conditionals'. *Foundations of Language* **6**, 89–94.

Adams, E.: 1975, *The Logic of Conditionals*. D. Reidel.

Anderson, A. R. and N. Belnap: 1975, *Entailment: The Logic of Relevance and Necessity*. Princeton University Press.

Appiah, A.: 1984, 'Generalising the Probabilistic Semantics of Conditionals'. *Journal of Philosophical Logic* **13**, 351–372.

Asher, N.: 1995, 'Commonsense Entailment: A Conditional Logic for Some Generics'. In: G. Crocco, L. F. del Cerro, and A. Herzog (eds.): *Conditionals: From Philosophy to Computer Science*. Clarendon Press, pp. 103–146.

Austin, J.: 1956, 'Ifs and Cans'. *Proceedings of the British Academy* **42**, 109–132.

Barwise, J. and R. Cooper: 1980, 'Generalized Quantifiers and Natural Language'. *Linguistics and Philosophy* **4**, 159–219.

Bennett, J.: 1974, 'Counterfactuals and Possible Worlds'. *Canadian Journal of Philosophy* **4**, 381–402.

Bennett, J.: 1988, 'Farewell to the Phlogiston Theory of Conditionals'. *Mind* **97**, 509–527.

Bennett, J.: 2003, *A Philosophical Guide to Conditionals*. Clarendon Press.

Bonevac, D., J. Dever, and D. Sosa: 2006, 'The Conditional Fallacy'. *The Philosophical Review* **115(3)**, 273–316.

Dancygier, B. and E. Mioduszewska: 1984, 'Semanto-Pragmatic Classification of Conditionals'. *Studia Anglica Posnaniensia* **17**, 121–133.

Dudman, V.: 1991, 'Interpretations of If-Sentences'. In: F. Jackson (ed.): *Conditionals*. Oxford University Press.

Edgington, D.: 1995, 'On Conditionals'. *Mind* **104**, 235–329.

Fine, K.: 1975, 'Review of Lewis's *Counterfactuals*'. *Mind* **84**, 451–458.

Gentzen, G.: 1969, 'Investigations Into Logical Deduction'. In: M. Szabo (ed.): *Collected Papers of Gerhard Gentzen*. North-Holland.

Gibbard, A.: 1981, 'Two Recent Theories of Conditionals'. In: W. Harper (ed.): *Ifs*. Dordrecht.

Gillies, A.: 2001, 'A New Solution to Moore's Paradox'. *Philosophical Studies* **105**, 237–250.

Gillies, A.: 2006, 'Counterfactual Scorekeeping'. ms, University of Michigan.

Gillies, A.: 2007, 'Counterfactual Scorekeeping'. *Linguistics and Philosophy* **30**, 329–360.

Girard, J.-Y.: 1987, 'Linear Logic'. *Theoretical Computer Science* **50**, 1–102.

Goodman, N.: 1954, *Fact, Fiction, and Forecast*. Harvard University Press.

Harman, G.: 1986, *Change in View*. The MIT Press.
Humberstone, L.: 2008, 'Conditionals'. ms, Monash University.
Jackson, F.: 1987, *Conditionals*. Basil Blackwell.
Kaufmann, S.: 2004, 'Conditioning Against The Grain: Abduction and Indicative Conditionals'. *Journal of Philosophical Logic* **33**, 583–606.
Krazter, A.: 1989, 'An Investigation of the Lumps of Thought'. *Linguistics and Philosophy* **12**, 607–653.
Kripke, S.: 1965, 'Semantical Analysis of Modal Logic, II: Non-Normal Modal Propositional Calculi'. In: L. H. J.W. Addison and A. Tarski (eds.): *The Theory of Models*. North-Holland.
Lapierre, S.: 1995, 'Conditionals and Quantifiers'. In: J. van der Does and J. van Eijck (eds.): *Quantifiers, Logic, and Language*. Stanford University Press.
Lewis, D.: 1973, *Counterfactuals*. Harvard University Press.
Lewis, D.: 1975, 'Adverbs of Quantification'. In: E. Keenan (ed.): *Formal Semantics of Natural Language*. Cambridge University Press.
Lewis, D.: 1979, 'Counterfactual Dependence and Time's Arrow'. *Nous* **13**, 455–476.
Morreau, M.: 1997, 'Fainthearted Conditionals'. *Journal of Philosophy* **94**, 187–211.
Ramsey, F.: 1931, 'General Propositions and Causality'. In: R. Braithwaite (ed.): *Foundations of Mathematics*. Routledge.
Sadock, J.: 1977, 'Modus Brevis: The Truncated Argument'. *Chicago Linguistic Society* **13**, 545–554.
Scott, D.: 1970, 'Advice on Modal Logic'. In: K. Lambert (ed.): *Philosophical Problems in Logic*. Reidel.
Segerberg, K.: 1971, 'An Essay in Classical Modal Logic'. *Filosofiska Studier* **13**.
Sher, G.: 1991, *The Bounds of Logic: A Generalized Viewpoint*. The MIT Press.
Siegel, M.: 2006, 'Biscuit Conditionals: Quantification Over Potential Literal Acts'. *Linguistics and Philosophy* **29**, 167–203.
Stalnaker, R.: 1968, 'A Theory of Conditionals'. In: *Studies in Logical Theory*. Basil Blackwell.
Thomason, R. and A. Gupta: 1980, 'A Theory of Conditionals in the Context of Branching Time'. *The Philosophical Review* **89**, 65–90.
van Benthem, J.: 1984, 'Foundations of Conditional Logic'. *Journal of Philosophical Logic* **13**, 303–349.
van Benthem, J.: 1986, *Essays in Logical Semantics*. D. Reidel Publishing Company.
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*. MIT Press.
von Fintel, K. and S. Iatridou: 2004, 'What to Do If You Want to Go to Harlem: Anankastic Conditionals and Related Matters'. ms, MIT.
Weatherson, B.: 2008, 'Conditionals and Indexical Relativism'. ms, Rutgers University.