

TRUTH AND CONSEQUENCE:
TOWARD A NEW ACCOUNT OF INDICATIVE CONDITIONALS

by

KELLY MARIE WEIRICH

B.A. Whitworth University, 2009

M.A. University of Colorado, 2012

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirement for the degree of
Doctor of Philosophy
Department of Philosophy

2015

This thesis entitled:
Truth and Consequence: Toward a New Account of Indicative Conditionals
written by Kelly Marie Weirich
has been approved for the Department of Philosophy

(Prof. David Barnett)

(Prof. Eric Chwang)

Date_____

The final copy of this thesis has been examined by the signatories, and we
find that both the content and the form meet acceptable presentation standards
of scholarly work in the above mentioned discipline.

Weirich, Kelly Marie (Ph.D., Philosophy)
Truth and Consequence: Toward a New Account of Indicative Conditionals
Thesis directed by Associate Professor David Barnett

Conditionals—sentences of the form ‘If A, B’—are ubiquitous in human discourse and reasoning, and yet giving a rigorous account of their meaning, or what it takes for true conditionals to be true, has proven difficult. There are three major accounts of conditionals in philosophy of language, representing widespread disagreement. In this dissertation, I explore answers to the following questions with respect to a class of conditionals called indicative conditionals: *What do conditionals mean? Are conditionals the sorts of things that are true or false? If so, what does it take for a true conditional to be true?* I present challenges for the major accounts and propose a new account called the Consequence Account. In Chapter One, I draw on linguistics to argue that a popular way of defending the Material Implication Account should be put to rest. In Chapter Two, I present a test for a good account of meaning, called the Relevance Test, and argue that the Material Implication Account fails this test. In Chapter Three, I present a challenge for any account according to which all conditionals with a true antecedent and a true consequent are true, which challenges versions of all three of the major accounts. In Chapter Four, I turn from arguing against the major accounts to arguing for a new kind of account. In this chapter, I defend the claim that conditionals have truth values against popular arguments to the contrary, and I present independent reasons for ascribing truth values to indicative conditionals. Finally, in Chapter Five, I present the Consequence Account, my own truth-valued account according to which an indicative conditional ‘If A, B’ is true just in case there is some relation by virtue of which B is a consequence of A. This account draws on the perception that a conditional expresses a certain kind of strong connection between the antecedent and consequent. In all of this, I hope to establish that the Consequence Account succeeds in areas in which the major accounts face challenges, and that it deserves consideration alongside them as a compelling, intuitive account of indicative conditionals.

For Scott, who loves me omniconditionally—
which is to say, unconditionally.

ACKNOWLEDGEMENTS

I would not have been able to write this dissertation without the support, time, expertise, and encouragement of many. I would like to thank the many people who helped me complete and refine this dissertation and the institutions that funded my research. First, thanks to my parents, who no doubt taught me my first conditionals and have supported me unconditionally—even when I declared a philosophy major. Thanks to David Barnett for being my advocate and critic throughout this whole process, and for letting me disagree with him on almost everything. Thanks to the many faculty who gave me helpful feedback on parts of this dissertation, including Graeme Forbes, Graham Oddie, Mitzi Lee, and Wes Morrison. Thanks to those who attended my colloquium at CU in fall of 2014 for energizing feedback on the parts of my dissertation presented there, and to the Philosophy Department for inviting me to speak. Thanks to my dissertation reading group—Rebecca Chan, Ben Rohrs, and Alex Zambrano—for reading drafts. Thanks to Professor Susan Felch of Calvin College for mentoring me through graduate school and giving an outside perspective on my profession. Finally, thanks to my husband Scott, who knows what it is like to write a dissertation. You have no idea how much that has meant.

Many institutions within the University of Colorado have supported my dissertation writing as well. Thanks to the Graduate School for funding me over the summer of 2013, to the CU Philosophy Department for a semester-long dissertation fellowship in spring of 2014, and to the Center for Humanities and the Arts for funding me for the academic year of 2014-2015 via the Thomas Edwin Devaney Dissertation Fellowship.

CONTENTS

INTRODUCTION: SCOPE, IMPORT, AND METHODOLOGY	1
I. Scope	2
A. What indicative conditionals are	2
B. Standard vs. non-standard conditionals	4
II. The major accounts	6
III. Methodology: on taking speakers' judgments seriously	7
IV. What's to come	8
CHAPTER ONE: MATERIAL IMPLICATION AND MERE IMPLICATURE	10
I. Material Implication and the Gricean Apparatus	11
II. Four failures of conversational implicature	15
A. Feature One: that the conventional meaning of the utterance be known	16
B. Feature Two: that the purported implicature be defeasible/cancelable	18
C. Feature Three: that the purported implicature be freely reinforceable	23
D. Feature Four: that the purported implicature be non-detachable	25
III. Conclusion	27
CHAPTER TWO: THE RELEVANCE TEST	29
I. Two main kinds of objection to the Material Implication Account	30
II. The Relevance Test	32
III. How the MIA fails the Relevance Test	36
IV. Objection	41
V. Conclusion	42

CHAPTER THREE: TT CONDITIONALS	43
I. How the major accounts fail to allow for false or unacceptable TT conditionals	44
A. The Material Implication Account	45
B. Possible Worlds Accounts	45
C. Suppositional/Probabilistic Accounts	46
II. In need of flexibility	51
A. The first phenomenon	51
1. False TT conditionals whose component clauses are irrelevant to each other	52
2. False TT conditionals that do not include irrelevance	59
3. A note on ‘whether or not’	64
B. The second phenomenon	66
III. How some accounts try to accommodate (P1) and (P2)	69
A. Take refuge in the “Assert the Stronger” principle.	69
B. Select for multiple worlds.	71
C. Distance the account from the actual world.	73
IV. Conclusion	75
CHAPTER FOUR: IN FAVOR OF TRUTH VALUES	76
I. Initial reasons for thinking conditionals have truth values	77
A. The speech acts in which we use conditionals	77
B. The use of conditionals in logic	78
C. Embedding	79
II. Argument One: Lewis’s Triviality Result and the Role of Conditional Probabilities	80
A. Lewis’s Triviality Result	80
B. Rejecting Lewis’s Assumption	83
C. Why Lewis’s Assumption was so initially plausible and where it fits in a good account	87
III. Argument Two: Edgington’s Dilemma	89
A. Edgington on TT Conditionals	90
B. Edgington on TF Conditionals	91
C. Edgington on FT Conditionals	92

D. Edgington on FF Conditionals	93
IV. Argument Three: Gibbard's Sly Pete Case	95
A. Gibbard's Sly Pete Case	95
B. Zack was wrong: responding to Gibbard	97
V. Argument Four: Barnett's Conscious Being Case	99
VI. The benefits of a "worldly" view	102
VII. Conclusion	105
CHAPTER FIVE: THE CONSEQUENCE ACCOUNT	107
I. The Consequence Account and the logical properties of 'if'	110
A. What's in and what's out	112
1. Unrelated events	112
2. P's raising the probability of Q	114
3. P's raising the probability of Q to some high degree	115
4. Purely epistemic associations	116
5. Backwards evidential conditionals	119
6. Accidental universal generalizations	122
B. Valid and invalid inferences	124
1. Modus ponens and modus tollens	124
2. Antecedent strengthening	125
3. Transitivity	126
II. What the Consequence Account does right	127
A. The Consequence Account has intuitive appeal.	127
B. The Consequence Account assigns truth values to conditionals.	128
C. There is no need for indeterminate truth values for conditional predictions whose antecedents do not obtain.	128
D. The Consequence Account explains the ambivalence of certain cases.	131
E. The Consequence Account makes the right judgments about 'TT' conditionals.	132
F. The Consequence Account gets the Sly Pete case right.	135
G. The Consequence Account deals with Barnett's Conscious Being Case.	136
H. The Consequence Account does not suffer from Gallimore's Problem.	138
I. The Consequence Account explains certain peripheral uses of conditionals.	139
J. The Consequence Account respects relevance.	141

K. The Consequence Account does not have to tell a separate story about assertability.	142
III. Objections and Replies	142
A. Grice's objection to connection-type accounts	142
B. Mackie's argument that connection-type accounts are circular	145
C. Edgington's argument against connection-type accounts	147
IV. Conclusion	149
CONCLUSION: TRUTH AND CONSEQUENCE	150
BIBLIOGRAPHY	152
APPENDIX	157

INTRODUCTION

SCOPE, IMPORT, AND METHODOLOGY

“Your ‘if’ is the only peacemaker: much virtue in ‘if.’”
—Touchstone, *As You Like It* (5.4.105-106)

‘If’ may be the most provocative two-letter word in English. ‘If’ can initiate a flight of fancy: *If I found a magic ring, I would use it only for good. If you give a mouse a cookie, he’s going to ask for a glass of milk.*¹ We find it in scientific reasoning: *If I add the acid, the protein will unfold.* We find it in politics: *If they pass this bill, the dollar will drop.* We find it in the most mundane of sentences: *If the glass falls, it will break. If I go to the store, I will pick up some coffee. If I keep writing, I will finish this dissertation.*

All of the above examples are conditionals: sentences of the form ‘If A, B’. Conditionals allow us to connect our ideas in a general way. They allow us to suppose something without committing to it. Such an enterprise is central to reasoning itself, and thus to philosophy as a discipline. Where would the philosopher be if she couldn’t suppose a view in order to see where it leads, or if her language had no means for expressing this reasoning? As hinted by Touchstone in Shakespeare’s *As You Like It*, conditionals allow people who disagree to suspend or disguise their disagreement in order to continue conversing. Conditionals are ubiquitous in all areas of human discourse, and we all seem to know what they mean. (When my niece was four, she understood ‘If you don’t finish your dinner, you won’t get dessert’.)

However, giving an account of the meaning of conditionals, and understanding what makes true conditionals true, is more difficult than one might expect. Conditionals are unlike other

¹ This sentence is the opening line from Laura Numeroff’s (1985) delightful children’s book *If You Give a Mouse a Cookie*.

locutions that take whole clauses as parts, for which there are straightforward truth conditions.

There is nothing similarly puzzling about ‘A and B’ or ‘A or B’. The conditional is distinctive in lacking wide agreement on its meaning or truth conditions, if any. This is a classical problem in the philosophy of language, and there is wide disagreement about the right account of conditionals.

In this dissertation, I explore answers to the following questions with respect to a class of conditionals called indicative conditionals: *What do conditionals mean? Are conditionals the sorts of things that are true or false? If so, what does it take for a true conditional to be true?* I present challenges for the major accounts in the philosophical literature and propose a new account called the Consequence Account. This introduction briefly presents the scope of my project, the major accounts of conditionals with which I am interacting, a word about methodology, and an outline of what is to come in the rest of the dissertation.

I. Scope

A. What indicative conditionals are

There are various kinds of conditionals, and the broadest distinction among them is between indicative and subjunctive conditionals. The line between indicative and subjunctive conditionals has been drawn in different places, but I will make the following simple grammatical distinction: a subjunctive conditional includes the helping verbs ‘would’, ‘were’ (plus infinitive), ‘has’, ‘had’, or ‘have’, and an indicative conditional does not. Indicative and subjunctive conditionals seem to have different meanings. Consider the following famous example (due to Adams (1970)):

- (1) If Oswald didn’t kill Kennedy, someone else did. [indicative]
- (2) If Oswald *hadn’t* killed Kennedy, someone else *would* have. [subjunctive]

The conditionals (1) and (2) seem to mean different things, and they seem to have different truth conditions. The distinction between conditionals such as (1) and (2) motivates the treatment of indicative and subjunctive conditionals as possibly requiring different accounts. It is worth noting that

some have argued that what I treat as future-tense indicatives have the same truth conditions as subjunctives. ‘If you strike the match, it will light’ seems equivalent to ‘If you were to strike the match, it would light’. (See, for example, Dudman (1988, 1991) and Bennett (1988).) Some examples that I discuss in Chapter Three undermine this supposed equivalence. For example, ‘If Sophia rolls the die, the die will come up six’ seems true just in case Sophia rolls a six, whereas ‘If Sophia were to roll the die, the die would come up six’ seems true only if the die is not fair—that is, only if something guarantees a six, given a roll of the die. In general for future indicative conditionals of this kind, it may be that some outcome *will* come about when it is not the case that that outcome *would* come about (or *would have* come about), given the circumstances. Although this consideration does not settle the matter, I will for this reason treat future-tense conditionals lacking the hallmark auxiliary verbs of subjunctive conditionals as indicatives. (For agreement, see, for example, Jackson (1990) and Bennett (1995).) While I do not argue that indicative and subjunctive conditionals cannot be accommodated with a single account, I will restrict the scope of my dissertation to indicative conditionals alone. Should it end up illuminating some aspect of subjunctive conditionals, so much the better.² Hereafter, the word ‘conditional’ on its own should be understood to mean ‘indicative conditional’, unless otherwise specified.

Conditionals plausibly also come in multiple forms besides ‘If A, B’. For example, ‘You won’t win unless you buy a ticket’ seems equivalent to ‘If you don’t buy a ticket, you won’t win’. A conjunction can also seem to be equivalent to a conditional. For example, ‘Move and you’ll blow the place to bits’ seems equivalent to ‘If you move, you’ll blow the place to bits’. There are many additional examples of sentences that seem equivalent to conditionals, but in the interest of streamlining the discussion, I will deal only with the standard ‘If A, B’ form. I will refer throughout

² For further discussion of the controversy on how to distinguish indicative and subjunctive conditionals and whether or not they differ importantly, see Edgington (1995), pp. 236-240, as well as Ellis (1978) and (1984). See Dudman (1983, 1988) for a distinction that cuts across grammatical subjunctive/indicative lines.

to A as the antecedent and B as the consequent, omitting single quotes around the letters representing those clauses for ease.

B. Standard vs. non-standard conditionals

While there are seemingly different ways to express a conditional, there are two kinds of conditionals that I want to flag as non-standard and thus as not expected to fall under an account of a typical conditional, though some accounts do not make these distinctions. Both of these locutions differ importantly from the kind of conditional found in the examples at the beginning of this introduction, which I will be treating as instances of standard conditionals. When I refer to some conditionals as “standard” and some as “non-standard,” I do not mean to imply that there is something less appropriate about the uses of certain kinds of conditionals. An utterance of these kinds of conditional may be perfectly germane. Nor do I intend to suggest that non-standard conditionals are fully idiomatic or otherwise unrelated to the central use of conditionals. The standard conditionals on which I focus are treated as central because they are, in my experience, statistically normal, and the philosophical literature on them is interesting. The conditionals that I will treat as non-standard are biscuit conditionals (and speech-act conditionals more generally) and semifactuals.

Biscuit conditionals are conditionals such as J.L. Austin’s (1961) ‘There are biscuits on the sideboard if you want some’. This sentence, while it can be rearranged in the form ‘If A, B’, does not seem to have the same kind of meaning as the standard conditionals at the beginning of this introduction. A biscuit conditional seems to extend a kind of offer, and—contrary to standard conditionals—it seems to presuppose the truth of its consequent. Biscuit conditionals are a subcategory of linguist Eve Sweetser’s (1991) category of speech-act conditionals: conditionals in which the ‘if’-clause indicates the circumstances that lead the speaker to make the speech act—for

example, ‘If you need help, my name is Kelly’.³ In such conditionals, the truth of the consequent is neither understood nor presented as conditional on the truth of the antecedent, and these conditionals likewise presuppose the truth of their consequents in cases in which the consequents have truth values (excluding, for example, cases in which the consequent is an imperative, as in ‘If you need anything, let me know’). For these reasons, biscuit conditionals and other speech-act conditionals are non-standard.

The second kind of non-standard conditional is what has been variously referred to as a semifactual or non-interference conditional. (I prefer ‘semifactual’.) Semifactuals are conditionals that either include or readily admit the addition of ‘even’ or ‘still’—or, to borrow from Douven (2008), conditionals whose acceptability/assertability rises with the addition of these words. For example, ‘[Even] if I am sick, I will [still] attend the ceremony’. Many philosophers⁴ consider semifactuals as non-standard conditionals, for good reason: semifactuals seem to assert their consequents, and standard conditionals do not.

To say that speech-act conditionals and semifactuals are non-standard is not to claim that they need not be accounted for in a complete grammar, nor to claim that their meanings are wholly unrelated to the meanings of standard conditionals. Rather, it is to claim that an account of conditionals should not be judged to fail if it does not treat these kinds of conditionals in the same way as standard conditionals. In fact, given their important differences, it would be surprising if an account of conditionals successfully treated speech-act conditionals and semifactuals exactly the same as it treats standard conditionals. The aforementioned features of speech-act conditionals and semifactuals give us good reasons for thinking that they deserve different treatment from standard

³ What I call standard conditionals correspond best to Sweetser’s “content conditionals,” though I do not want to rule out at the outset that one can provide a unified theory of both content conditionals and what Sweetser refers to as “epistemic conditionals.” I return to this question in chapter five.

⁴ Among them are Douven (2008, pp. 31-32), Burgess (2004, p. 567), and Lycan (2001, pp. 30-31).

conditionals, and far from leading to an objectionably disunified account, this different treatment is exactly what one should expect.

Any worries about disunity aside, one might worry that the treatment of certain kinds of sentences of the form ‘If A, B’ as non-standard violates compositionality in grammar. After all, this treatment entails that there are multiple kinds of meaning for syntactically isomorphic sentences. However, there are grammars that treat locutions as falling on a continuum from the wholly compositional to the wholly idiomatic—rather than being strictly one or the other—and such an approach can helpfully categorize the conditionals I deem to be non-standard as somewhat compositional. For an explanation of how such a treatment of conditionals need not violate the spirit of a compositional grammar, see Kay and Michaelis (2012).

II. The major accounts

Throughout, I will refer to three kinds of account of indicative conditionals as the major accounts. These are the Material Implication Account, the Possible Worlds Account, and Suppositional/Probabilistic Accounts. These three accounts have dominated the literature on indicative conditionals in the last fifty years. Briefly, the Material Implication Account is the account according to ‘If A, B’ in natural language has the same truth conditions as ‘ $A \supset B$ ’ in classical logic—in other words, ‘If A, B’ is false when A is true and B is false, and otherwise it is true. The Material Implication Account’s truth conditions are often supplemented with a story about what makes certain conditionals assertable and others unassertable. Proponents include Grice (1989), Jackson (1979, 1981, 1987), Lewis (1976), and Mellor (1993).

The Possible Worlds Account holds that ‘If A, B’ is true at world w if and only if B is true in the closest (set of) A-world(s) to w . (I will follow convention by using ‘A-world’ as shorthand for ‘possible world at which A is true,’ as well as similar variants.) Closeness of worlds is a matter of

similarity. Stalnaker (1968, 1970) set out the canonical Possible Worlds Account for indicative conditionals. Other proponents include Ellis (1984), Weatherson (2001), and Nolan (2003).

There are a few different accounts that fall under the label of Suppositional/Probabilistic Accounts. Most of these accounts hold that a conditional ‘If A, B’ has a probability equal to the conditional probability of B, given A, or else that the assertability of the conditional goes by this conditional probability. See, for example, Adams (1965, 1966, 1975), Mackie (1973), and Edgington (1986, 1995). The version by David Barnett (2006) holds that uttering the statement ‘If A, B’ is an act expressing a proposition whose content is B and whose context is the supposition that A.

One might wonder why the work of relevance logicians such as C.I. Lewis and Graham Priest does not appear on the list of major accounts. The reason is this: the work on relevance logic, for better or worse, has developed as something of a separate literature from the literature dealing with the major accounts listed above. At any rate, those with whose work I am interacting do not reference it at all. Thus, I restrict my discussions to the three major accounts above and their variants.

III. Methodology: on taking speakers’ judgments seriously

Before the body of the dissertation, I want to make a brief point about the methodology I employ. While there is reasonable caution in ascribing too much philosophical weight to people’s unreflective judgments and practices, philosophy of language is one sub-discipline in which the judgments and practices of competent speakers deserve a great deal of consideration. Languages are somewhat special in this respect: what they are depends on what they are intended to be. Since these intentions and the rules they engender are not always made explicit, one must reason about them in order to come up with a consistent and accurate account of the parts of a language. An account of indicative conditionals, if it is to reach its target, must glide along the twin rails of what is possible

for a language—what is reasonable, what is internally consistent—and what is intended by a language—what it is that people set out to do when they think, speak, or write in it. Since I have good reasons to believe that I am a competent user of English (my Bachelor of Arts in English Writing being among them), my methodology involves presenting my own judgments about particular uses of conditionals and nudging at the reader to share my judgment, as is common in this area. When dealing with a natural language, the fact that an account is intuitive to competent speakers is among the highest points that can be in its favor. Thus, in what follows, I place a great deal of weight on fit with speakers' use and pre-theoretic judgments about the truth or acceptability of certain conditionals.

IV. What's to come

This dissertation begins with a series of independent chapters presenting original arguments against the major accounts (chapters One through Three). In Chapter One, I draw on linguistics to argue that a popular way of defending the Material Implication Account by means of Gricean conversational implicature should be put to rest, because the purported implicature does not have the features of conversational implicature. In Chapter Two, I present a test for a good account of meaning, called the Relevance Test, and argue that the Material Implication Account fails this test. In Chapter Three, I present a challenge for any account according to which all conditionals with a true antecedent and a true consequent (TT conditionals) are true. This argument challenges versions of all three of the major accounts. In Chapter Four, I make the turn from arguing against the major accounts to arguing for a new kind of account. In this chapter, I defend the claim that conditionals have truth values against popular arguments to the contrary, and I present independent reasons for ascribing truth values to indicative conditionals. Finally, in Chapter Five, I present the Consequence Account, my own truth-valued account according to which an indicative conditional 'If A, B' is true

just in case there is some relation by virtue of which B is a consequence of A. This account draws on the perception that a conditional expresses a certain kind of strong connection between the antecedent and consequent. In all of this, I hope to establish that the Consequence Account succeeds in areas in which the major accounts face challenges, and that it deserves consideration alongside them as a compelling, intuitive account of indicative conditionals.

CHAPTER ONE

MATERIAL IMPLICATION AND MERE IMPLICATURE

Abstract: The Material Implication Account of indicative conditionals holds that the indicative conditional ‘if P, then Q’ in English has the same meaning as the material conditional ‘ $P \supset Q$ ’ in classical logic. Well-known challenges to the Material Implication Account, known as the paradoxes of material implication, have led some of its proponents to appeal to conversational implicature to try to explain the paradoxical nature of the paradoxes of material implication. In this paper, I challenge such an appeal. First, I argue that a typical utterance of an indicative conditional communicates the existence of a certain kind of relation between the antecedent and consequent. Second, I argue that the Material Implication Account faces problems in accounting for the communication of such a relation, when measured against four of the characteristic features of mere conversational implicature. I conclude that versions of the Material Implication Account that appeal to conversational implicature are unsatisfactory.

The Material Implication Account of indicative conditionals holds that the indicative conditional ‘if P, then Q’ in English has the same meaning as the material conditional ‘ $P \supset Q$ ’ in classical logic—that is, that it is false when P is true and Q is false, and it is true otherwise. This account faces numerous objections of a similar nature, collectively referred to as the paradoxes of material implication. The paradoxes of material implication are cases in which what is permissible with respect to a material conditional is not permissible with respect to the corresponding indicative conditional: an inference involving a material conditional is permissible, but the inference involving the corresponding indicative conditional is not permissible; a material conditional is assertable in some context, but the corresponding indicative is not; and so on. In light of these so-called paradoxes,¹ many defenders of the Material Implication Account have adopted a Gricean apparatus

¹ As Davis (1979) says, “It is a paradox only to those who think material conditionals and natural language conditionals have the same truth conditions” (p. 550).

which, when added to the Material Implication Account, is meant to explain why there are apparent differences between the way material conditionals behave and the way indicative conditionals behave. I shall use ‘Gricean apparatus’ somewhat broadly as the term for a feature of a linguistic account according to which some communicable content—called an ‘implicature’—plays a role in the assertability of a phrase but does not factor into the truth conditions of that phrase or what is literally said by means of it. In this paper, I argue that any Gricean apparatus that is robust enough to explain the use of indicative conditionals cannot consist in mere conversational implicature. I would like to put to rest the thought that a version of the Material Implication Account can account for purported implicatures as being merely conversational.

In section I, I catalogue the versions of the Material Implication Account that add such a Gricean apparatus and explain what role this apparatus is supposed to play. I present a version of Strawson’s (1981) suggestion that, in indicative conditionals, the antecedent is being put forward as a ground for the consequent. I argue that any Gricean apparatus must account for the communication of this relation in order to present an adequate account of indicative conditionals. Then, in section II, I explain how the apparatus of conversational implicature fails to account for the communication of this relation. Four characteristics of mere conversational implicature fit ill with the way indicatives behave with respect to the communication of this relation. Finally, in section III, I discuss the further implications this analysis has for the defensibility of the Material Implication Account.

I. Material Implication and the Gricean Apparatus

The Material Implication Account of indicative conditionals faces well-known challenges.² The so-called paradoxes that plague it are too numerous to list here helpfully, but here are a few examples:

² See, for example, Stalnaker (1975), pp. 136-137, and Davis (1979), pp. 550-551.

(A) ‘ $\sim P \therefore P \supset Q$ ’ is valid and assertable, but the following is invalid and unassertable: ‘She did not eat the apple. Therefore, if she did eat the apple, she was ill.’³

(B) ‘ $Q \therefore P \supset Q$ ’ is valid and assertable, but the following is invalid and unassertable: ‘The earth will not be plunged into darkness in eighteen minutes. Therefore, if the sun goes out of existence in ten minutes, the earth will not be plunged into darkness in eighteen minutes.’⁴

(C) ‘ $P \vee \sim P \therefore (Q \supset P) \vee (P \supset R)$ ’ is valid and assertable, but the following is invalid and unassertable: ‘Either dinosaurs existed or it is not the case that dinosaurs existed. Therefore, either if the fossil record is wildly misleading, dinosaurs existed, or else if dinosaurs existed, all dinosaurs were carnivores.’

These paradoxes suggest that ‘if’ in English is not equivalent to ‘ \supset ’ in classical logic. In light of these and other paradoxes of material implication, proponents of the Material Implication Account have often supplemented the account with a Gricean apparatus, meant to explain the seemingly paradoxical nature of the inferences. Again, a ‘Gricean apparatus’ is a feature of a linguistic account according to which some communicable content—called an ‘implicature’—plays a role in the assertability of a phrase but does not factor into its truth conditions or what is literally said by means of that phrase. The implicature involved in a Gricean apparatus may take one of two forms: conversational or conventional. Conversational implicature arises as a result of assuming that one’s interlocutor is cooperating with certain understood rules for linguistic utterance, referred to as conversational maxims.⁵ Conventional implicature, on the other hand, is a phenomenon in which what is suggested arises not from assumptions of cooperation with conversational maxims, but rather from the non-truth-conditional meaning of the utterance. While the line between pragmatics and semantics is sometimes hard to draw,⁶ conversational implicature is closer to the pragmatic end,

³ See Edgington (1995), p. 243.

⁴ See Jackson (1979), p. 114.

⁵ These rules are the maxim of Quantity (“Make your contribution as informative as is required [for the current purposes of the exchange]. Do not make your contribution more informative than is required.”), the maxim of Quality (“Try to make your contribution one that is true.”), the maxim of Relation (“Be relevant.”), and the maxim of Manner (“Be perspicuous.”). These four maxims are unified under Grice’s Cooperative Principle: “Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.” (Grice (1989), pp. 26-27)

⁶ See, for example, Levinson (1983) chapter one and Saeed (2009), pp. 17-18.

whereas conventional implicature is closer to the semantic end. For this reason, it seems appropriate to consider conversational implicature to be *mere* implicature, by way of contrast with conventional implicature, which is in some sense part of the meaning of the words uttered.

Proponents of the Material Implication Account whose Gricean apparatus is meant to be conversational, as opposed to conventional, include Clark (1971, 1974), Quine (1974), and Grice (1989). Clark (1971), responding to Strawson's (1981) account of what is meant by the conditional, says, "The implications may arise, not from the meaning of 'if,' but in virtue of some general conversational requirements or conventions about relevance, point, etc." (p. 34). Though Clark uses the term 'conventions', it is clear that Clark is referring to Grice's conversational maxims, which factor into conversational—and not conventional—implicature. According to Quine (1974), some kind of relevance between the antecedent and consequent is merely a requirement for the conditional to be usefully stated, and does not enter into its meaning. Finally, Grice (1989) himself thought the proper account of the conditional was a Material Implication Account, and that the paradoxes of material implication should be accounted for in terms of what Jackson (1979) later referred to as the "assert the stronger" principle. One should not utter 'if P, then Q' on the basis of belief in either the falsity of P or the truth of Q, because in that case one should utter the logically stronger 'not P' or 'Q'. Those who add a Gricean apparatus that is conventional, rather than conversational, include Jackson (1979) and Lewis (1976). Jackson holds that, while the truth conditions of 'if P, then Q' in English are the same as those of $P \supset Q$ in classical logic, uttering the conditional in English has the conventional implicature that such an utterance is robust with respect to P—that is to say, that learning P would not render 'if P, then Q' improbable. While Lewis initially endorsed a conversational Gricean apparatus, he announced a new endorsement of Jackson's view in the postscript to his (1976).

While the content of the Gricean apparatus varies from account to account, for the remainder of this paper I will focus on one particular kind of purported implicature, which is discussed by Strawson (1981): what Strawson calls the ‘ground-consequent’ relation. According to Strawson, ‘if’ is importantly like the coordinating conjunction ‘so’: it expresses a dependence relation without specifying the particular kind of dependence relation. Lycan (2001) likewise considers dependence between the consequent and antecedent to be one of “the most elementary and uncontroversial semantic and pragmatic features of paradigm conditionals” (p. 141). The common element among conditional utterances is that the antecedent is being advanced as a ground for the consequent. For example, in

(1) If you drop the glass, it will shatter.

your dropping the glass is presented as a rational ground for accepting that the glass will shatter, or perhaps as a causal ground for the glass’s shattering, should your dropping the glass obtain. The antecedent is presented as a reason for believing the consequent. I will refer to such a relation as the consequence relation. The communication of a consequence relation is central to our use of conditionals. This relation is arguably what Jackson (1979) tries to account for by means of his robustness condition: learning that the antecedent is true should not lead one to disbelieve the conditional, because learning that the antecedent is true has another conflicting function—namely, giving one grounds for believing the consequent. In fact, the suggestion of such a relation is so strong that some have presented accounts of indicative conditionals wherein this relation is central to their meaning.⁷ This relation forms the core of my own account, presented in Chapter Five. One reason for thinking that the consequence relation is central to our use of conditionals is that, as I argue in section II, its suggestion is non-defeasible (pace Grice (1989)). I leave that argument until section II. Here, I leave it open whether this relation is a purely epistemic one, such that the relata

⁷ For example, Mellor (1993) and Douven (2008).

are beliefs (as Mellor (1993) and Douven (2008) have it), or whether instead it is not primarily an epistemic relation, which allows for truth conditions that do not involve or reference a person's beliefs, as I will argue in Chapter Five.⁸ Any version of the Material Implication Account must account for the suggestion of a consequence relation between an indicative conditional's antecedent and consequent.

II. Four failures of conversational implicature

The consequence relation, which, as I argue above, must be accounted for by any account of indicative conditionals, is not well accounted for by the conversational implicature version of the Material Implication Account. Here, I present four characteristics of conversational implicature, present both in Grice (1989) and in subsequent linguistics literature, and argue that the Material Implication Account faces problems with respect to the consequence relation for all four characteristics. The four characteristics are

- (i) that the conventional meaning of the utterance be known,⁹
- (ii) that the purported implicature be defeasible/cancelable,¹⁰
- (iii) that the purported implicature be freely reinforceable,¹¹ and
- (iv) that the purported implicature be non-detachable.¹²

With the exception of (iii), this list occurs in Grice (1989), and all of (i)-(iv) occur in Levinson (1973, 2000). Since the Material Implication Account fails on all four counts with respect to the consequence relation, the suggestion of such a relation cannot consist in mere conversational

⁸ The purely epistemic version must respond to the question, "Good grounds for whom?" For example, I may well take the supposition that only one conscious being exists as grounds for believing that I myself am that being (since it's obvious to me that I exist and that I am a conscious being). I should not expect you to take the supposition that only one conscious being exists as grounds for believing that *I* am that being, assuming a self-conscious reader. Thanks to David Barnett for pointing out this difficulty.

⁹ Levinson (1983), p. 113; Grice (1989), p. 49; Levinson (2000), p. 15.

¹⁰ Levinson (1983), p. 114; Grice (1989), p. 44; Levinson (2000), p. 15.

¹¹ Levinson (1983), p. 120 (with reference to Sadock (1978)); Levinson (2000), p. 15.

¹² Levinson (1983), pp. 116-117; Grice (1989), p. 39; Levinson (2000), p. 15.

implicature. In the cases of the first three features, I present my own challenges, and in the case of the fourth feature, I appeal to a challenge deriving from Bach (1999).

The discussion of these challenges involves appeals to intuitive judgments about whether or not a certain phenomenon, such as defeasibility, occurs with respect to a particular utterance. Judgments in some cases may differ, but even if not all of the challenges are equally compelling, together they comprise a total case that tells strongly against the conversational implicature version of the Material Implication Account. Given the deeply counterintuitive consequences of the Material Implication Account, this account can only be as successful as the semantic or pragmatic story that buoys it. Other well-developed alternatives to Grice's theory of pragmatics have been given,¹³ but here I would like to put to rest the optimism I have encountered—more often in person than in print—that a sort of vaguely Gricean picture should do the trick. A picture that is even somewhat faithful to Grice's will be deeply problematic, as I show in the following sub-sections. At almost every turn, conversational implicature fails to account for the communication of the consequence relation. If the Material Implication Account is true, some pragmatic story or other must be told; but the conversational version cannot be the right one.

A. Feature One: that the conventional meaning of the utterance be known

Grice (1989) says, in "Further Notes to Logic and Conversation,"

... If nonconventional [i.e., conversational] implicature is built on what is said, if what is said is closely related to the conventional force of the words used, and if the presence of the implicature depends on the intentions of the speaker, or at least on his assumptions, with regard to the possibility of the nature of the implicature being worked out, then it would appear that the speaker must (in some sense or other of the word *know*) know what is the conventional force of the words he is using. (49)

¹³ See, for example, relevance theory as presented in Carston (2012).

Conversational implicature involves the assumption that the hearer will be able to reason out, by means of the conversational maxims, what is being implicated by uttering something with such-and-such conventional meaning within its context. (This is not to say that hearers must or do consciously reason out what is being implicated. The reasoning may be implicit, and the implicature may be so common as to be automatically understood. See Morgan's (1978) elucidation of short-circuited implicature for further discussion.) Thus, in order for conversational implicature to arise, the conventional meaning of the utterance must be known.

However, if the Material Implication Account is true, the conventional meaning of 'if P, then Q' is not known. Mackie (1973) agrees, saying, "[I]t is implausible to claim that English hearers generally do recognize 'if P as meaning primarily \supset ' (p. 80).¹⁴ Probably most people who have taken a course in sentential logic have had the experience of learning the truth table for ' $P \supset Q$ ' and protesting that this cannot possibly be what is meant by 'if P, then Q'. Logic books and some instructors construct examples to pump the intuition that 'if P and ' \supset ' have the same meaning. The reason for providing such intuition pumps is that it is not obvious that the meaning of 'if P is the same as the meaning of ' \supset '. Students need to be convinced. If they are convinced of the equivalence of the English construction with the construction in classical logic, they have learned something new, which they did not believe—much less know—prior to exposure with symbolic logic. This story is a familiar one. Even if it is true that 'if P, then Q' has the conventional meaning of ' $P \supset Q$ ', it is not plausible to suppose that the average competent speaker of English knows this to be the case. Of course, a speaker may in some sense implicitly know what is meant by an utterance without being able to articulate explicitly the meaning of that utterance. What is so striking in the case of the

¹⁴ Mackie goes on to note the possibility that Grice is not (or might not be) making a claim about people's actual psychology, but merely a claim that we can think of 'if P as if they were material conditionals, consistent with the data. But this assessment makes it hard to understand Grice's requirements about knowledge of the literal meaning of the phrase. This approach seems to do away with the knowledge requirement altogether.

Material Implication Account is that, once the equivalence of ‘if’ and ‘ \supset ’ is proposed for consideration, competent speakers of English reject it. I am not alone in pointing out this rejection. Jackson (1979, p. 115) mentions the fact that, though competent speakers often accept Grice’s explanation of ‘or’ as inclusive, they seldom accept Grice’s explanation of ‘if’ as ‘ \supset ’—and Bennett (2003, p. 33) affirms Jackson’s point. What is new about my criticism here is not the fact that competent speakers often reject a material-implication-plus-conversational-implicature account. The problem I point out is worse: that this very account *presupposes* that competent speakers already know (or, we may charitably weaken the claim to say: that speakers should accept when presented with) the conventional meaning of ‘if’ as ‘ \supset .’ Such is not the case. The Material Implication Account fails with respect to this first criterion, without our even needing to reference the consequence relation at all.

B. Feature Two: that the purported implicature be defeasible/cancelable

A purported implicature implicated by an utterance is *defeasible* just in case something further can be said in addition to the utterance that removes the implicature without rendering the expanded utterance infelicitous or self-contradictory. For example, if one is asked whether a particular student has a promising academic career, and one responds,

(2) He’s a snappy dresser, which couldn’t hurt.

one has conversationally implicated that the student does not have the more important qualities it takes to have a promising academic career. The utterer of (2) has said something irrelevant to the question asked, and so the hearer can be expected to work out the implication that the real answer to the question is not one that it would be polite to say out loud. However, such an implication can be defeated, for example, like this:

(3) He's a snappy dresser, which couldn't hurt. I also think he has the brains and the work ethic to do very well for himself in academia.

By adding to the original utterance, the speaker has nullified the suggestion that the student lacks the qualities that make for a good academic. Following Levinson (1983), I prefer 'defeasible' to 'cancelable,' since one may remove a purported implicature without denying the truth of what is potentially implicated.¹⁵ For example, the purported implicature generated by (2) could be defeated by (4) without the speaker's denying its truth.

(4) He's a snappy dresser, which couldn't hurt. But I haven't had any interactions with him in a professional setting, so I can't speak to his real qualifications.

Defeasible implicatures are contrasted with implicatures or entailments that cannot be removed without rendering the expanded utterance infelicitous or self-contradictory. Consider (5)

(5) Lane is poor, but she is famous.

which implicates (6)

(6) There is some contrast between being poor and being famous.

Such implicature cannot be defeated. For (7) comes off as odd and self-defeating.

(7) Lane is poor, but she is famous—though I do not mean to suggest that there is a contrast between being poor and being famous.

The utterer of (7) can most charitably be understood to have retracted what is implicated by 'but,' for sincerely asserting all of (7) comes off as infelicitous.

One characteristic of conversational, as opposed to conventional, implicature is that mere conversational implicature is defeasible. However, pace Grice, the suggestion of a consequence

¹⁵ See p. 115.

relation between the antecedent and consequent is not defeasible.^{16,17} Instead of focusing on Grice's examples, I would like to try to construct my own case in which the suggestion of a consequence relation is defeated. In order to be a successful instance of defeat, the resulting utterance must not be either infelicitous or self-contradictory. It will be helpful for our purposes if the conditional is true or assertable according to the Material Implication Account (so as to grant as much to the Material Implication Account as possible), but is not an instance in which the antecedent is being presented as a good ground for believing the consequent (so as to open up the possibility for defeat).

Consider the following example. Suppose it is well known to both Mondo and Jarrell that Mondo favors Martinez as a candidate and that Martinez has no chance of winning the election.

Mondo: If Martinez wins the election, then she will ruin this city.

Jarrell: [quizzical look] I thought you favored Martinez....

Mondo: Oh, I do. I think she would do great work. I only said that because she clearly won't win the election. I didn't mean to suggest that her winning would lead to a ruined city.

Has Mondo defeated the implicature that Martinez's hypothetical winning is a good grounds for the belief that she will ruin the city? Charitably, Jarrell may assume that Mondo is being facetious (assuming the position of those incompetent voters who disfavor Martinez and mocking them merely by stating their position)—but that is a situation in which Mondo would be assumed not to be asserting the conditional at all. What we need is a situation in which Mondo *asserts* a conditional and the implication of a consequence relation is defeated. The present example is not one in which the purported conversational implicature is defeated, even though the speaker explicitly denies the

¹⁶ In his (1989), Grice presents examples of purportedly-defeated implicatures arising from conditionals (pp. 59-60). I admit that I find Grice's examples baffling and infelicitous, and I'm not alone. Adams (1975) says, "Even Grice's tests concerning what can be 'cancelled' may be argued both ways" (pp. 20-21). Since my response to Grice's cases is merely that the purported cancellations seem infelicitous, and judgments of infelicity are likely to differ somewhat from reader to reader, I direct the reader to Grice rather than rehearsing his cases here.

¹⁷ Douven (2007), p. 116, points out some problems with defeasibility for a particular series of statements, in response to Bradley's (2000) preservation condition, but does not argue that there are problems with defeasibility that render the conversational implicature version of the Material Implication Account untenable in general.

consequence relation *and* provides alternative reasons for stating the conditional (namely, the falsity of the antecedent). If the combination of these two elements is not enough to defeat the implication that there is a consequence relation between the antecedent and the consequent, it is difficult to see what could accomplish the task of defeating that implication.

Note that Mondo's conditional is not an instance of a Dutchman (or monkey's uncle) conditional—that is, a conditional that seems to deny its antecedent by means of an absurd (and absurdly unrelated) consequent, such as 'If that's true, I'm a Dutchman (monkey's uncle)'. (Dutchmen must, of course, use the monkey's uncle or other variants to achieve the same effect.) Dutchman conditionals are plausibly treated as somewhat non-standard because they seem to deny their antecedents, which conditionals in general do not do. Another reason for thinking that Dutchman conditionals are non-standard is that they seem to have what Bennett (1995) calls the Opt-out Property—that is, they are such that learning their antecedents would lead one to reject the whole conditional rather than to come to believe the consequent—which Bennett observes not to be a property of indicative conditionals. (Oddly, Mackie (1973, p. 84) thinks a connection-type account can accommodate Dutchman cases, since he thinks the metaphorical Dutchman is symbolic of "a stupid person, one lacking in judgment," in which case being wrong about the antecedent does have as a consequence that the speaker is one of those. This kind of case seems to me to be a distinct subset of these cases more generally.) Returning to Mondo's conditional, the idea that Martinez might ruin the city may not be probable to the participants of that conversation, but it is not like the kind of consequent we find in Dutchman conditionals. It is not absurd and absurdly unrelated to the antecedent, as the consequents of Dutchman conditionals are, and neither does it seem to deny its antecedent.

One of the difficulties in establishing non-defeasibility is that, if an implicature is non-defeasible, then it will be difficult to generate a charitable-sounding case in which defeat is

attempted. Assuming that the implicature really is non-defeasible, any attempt to defeat it will come off as clumsy, and that clumsiness may cause one to suspect that the author is not trying very hard to present a genuine instance of defeat. I think part of the difficulty may lie in this: sometimes one must explain oneself when defeating a potential implicature. For an explanation of defeated implicature where the potential for implicature is particularly strong, it is most natural to state not only that one is not implicating the potential implicature, but also the positive reason why one made the statement that carried the (now defeated) implication. In the case of Mondo above, it is difficult to see what could serve as the motivation for uttering the conditional in the absence of the consequence relation. There's no story that seems natural to tell, even allowing that we can add in whatever hypothetical features we want to Mondo's case. The very fact that one cannot easily imagine reasons for defeating the potential implicature of the consequence relation tells in favor of the claim that this implicature is non-defeasible. Defeating an implicature should be easy, especially for an audience that understands Grice's account. When it is difficult even to imagine the circumstances in which defeat would be appropriate, that itself is reason to believe that the purported implicature is non-defeasible.

One worry that has been raised for my argument in this section is that showing difficulty for defeasibility in one case does *not* show that *no* potential implicature of the consequence relation by the conditional can be defeated. However, defeasibility is a feature that instances of conversational implicature have in general, which means that, barring some very unusual context, any instance of conversational implicature should be defeasible. Showing defeasibility to be problematic or impossible in one case does not prove that it is impossible in every case, but it does suggest that defeasibility is not a feature that the conditional's communication of the consequence relation generally carries. We can usefully compare the difficulty of defeating the purported implicature of conditionals with the relative ease of defeating the purported implicature of disjunctions. Asserting a

disjunction carries the conversational implicature that the utterer believes only one disjunct to be true, though its truth conditions (according to most) are such that a disjunction is true just in case one *or more* disjunct is true. Yet the implicature in this case is easily defeated—for example, when one says, ‘You may have cake or you may have ice cream—or both’. The ease of defeating the implicature in the case of the disjunction stands in stark contrast with the difficulty in the case of conditionals. As Ellis (1984) points out, the comparison of ‘if’ with ‘or’ is especially poignant: “If [indicative conditionals] are material conditionals, as Jackson and Lewis hold, then they *are* disjunctions.” (p. 60). I submit that the purported implicature of a consequence relation is not, in general, defeasible, which presents a second divergence from the features of conversational implicature.

C. Feature Three: that the purported implicature be freely reinforceable

A purported implicature *x* generated by utterance *y* is freely reinforceable just in case *x* can be stated explicitly in addition to *y* without redundancy—that is, without repeating what is said. For example, consider utterances (8)-(10):

- (8) Has Daniel made dinner yet?
 (9) I think he’s still at work.

In this case, (9) conversationally implicates (10)

- (10) Daniel has not made dinner yet.

It is possible to state (10) explicitly, in addition to (9), in order to reinforce that it is implied.

- (11) I think he’s still at work. So, no, Daniel has not made dinner yet.

(11) does not involve redundancy, because its second sentence does not repeat what is said by its first sentence. Its first sentence says something about Daniel’s still being at work, and does not strictly say anything about whether or not Daniel has made dinner—though on its own it does suggest that Daniel has not made dinner yet. Since (11) does not involve redundancy, (9)’s

implicature of (10) is freely reinforceable, in keeping with the features of conversational implicature. Being freely reinforceable is a sort of counterpart to being defeasible. If a purported implicature is merely conversationally implied and is not generated more closely from the meaning of the words themselves as in conventional implicature, then it should be possible for one either to deny that this further implicature is intended or else to emphasize that it is, without self-contradiction in the first case or redundancy in the second.

Just as the purported implicature of the consequence relation is not defeasible, I submit that it is also not freely reinforceable. Attempts to reinforce it produce redundancy. The difference between a statement's being redundant and its not being redundant may be difficult to discern in some cases, but a good test seems to be to determine whether the statement sounds primarily *clarificatory*—in which case it is not redundant, because it adds something, namely clarification—or whether on the other hand it comes off as being *overly pedagogical*—in which case it is redundant, because the most charitable reading of it is as repeating something for the sake of the learner. This test is not as strong as some of the tests in other sections, because sometimes it is difficult to judge the difference between true redundancy and being explicit about something that has been communicated but not literally said. Nonetheless, this test give us a general guide in determining whether a purported implicature is freely reinforceable. The second sentence in example (11) above sounds primarily clarificatory. Now consider (12) and (13):

(12) If you drop the glass, it will shatter. Should you drop the glass, dropping the glass will have its shattering as a consequence.

(13) If Alison is in Boulder, then she is in Colorado. Alison's being in Boulder has her being in Colorado as a consequence.

In neither of these cases does the second sentence seem to express more than is expressed in the first. They do not communicate anything that is not already communicated in the first sentence.

While (12) and (13) might, in their entirety, be appropriate in a certain pedagogical setting, they have

that teacherly air of redundancy. The second sentence reinforces what is said in the first by restating it in other words, just as a teacher might say, for example,

(14) All electrons have negative charge. Every single electron—in your body, in this desk, in everything—has that same negative charge.

The problem is not that the attempt at reinforcement is infelicitous in all contexts, but that it is redundant. And so the second sentences in (12) and (13) are: they do not add anything that has not already been said. Thus, conditionals fail to exhibit yet another feature of mere conversational implicature with respect to the consequence relation. Such a relation, while communicated by the conditional, is not freely reinforceable as an instance of conversational implicature should be.

D. Feature Four: that the purported implicature be non-detachable

A purported implicature is detachable just in case it can be removed by restating the original utterance in truth-functionally-equivalent terms. Instances of conversational implicature are non-detachable. Stating them in truth-functionally-equivalent terms should, in the same context, produce an utterance that generates the same implicature.¹⁸ For example, consider again (2):

(2) He's a snappy dresser, which couldn't hurt.

(2) conversationally implicates, in the context introduced above, that the student in question lacks the more important qualities it takes to have a promising academic career. So does the truth-functionally equivalent (15):

(15) His attire is usually sharp, which will serve him well.

Stating the meaning of (2) in the equivalent terms of (15) does not result in detachment of the implicature; the implicature remains. If the consequence relation is merely conversationally

¹⁸ The prominent exception to this rule is the case of implicatures that arise as a result of application of the maxim of manner. Since manner consists in the style of delivery (brevity vs. prolixity, informal vs. formal register...), truth-functionally-equivalent terms cannot in general be expected to retain the features that generate the implicature. See, for example, Levinson (1983), p. 117.

implicated by a conditional, this implication should remain when the conditional is translated into truth-functionally-equivalent terms.

The problem the Material Implication Account faces with respect to non-detachability is not the same problem it faces with respect to the other three features, for it is not clear—to the author, anyway—whether or not purportedly truth-functionally-equivalent substitutions result in detachment. Nonetheless, as Bach (1999) argues, the test itself is flawed. As I will argue, the test of non-detachment at worst demonstrates that the implicature cannot be conversational and at best does nothing to rule out that the purported implicature is a part of the conventional meaning of the utterance.

The claim that a particular conditional is semantically equivalent to some other sentence will, in most cases, be controversial. However, since the current task is to argue against a certain version of Material Implication Account, it should be dialectically acceptable to use some sentence that the Material Implication Account holds to be equivalent to a conditional, even though this equivalence may be denied by those who do not endorse the Material Implication Account. In this way, we can address the Material Implication Account on its own terms. According to the Material Implication Account, ‘if P, then Q’ is truth-functionally equivalent to ‘not-P or Q,’ ‘Q or not-P,’ ‘not both-P-and-not-Q,’ and numerous other sentences. Let us consider the substitution of (16) with any of (17)-(19).

- (16) If you drop the glass, it will shatter.
- (17) Either you won’t drop the glass, or it will shatter.
- (18) Either the glass will shatter, or you won’t drop it.
- (19) It won’t be true that you both drop the glass and it doesn’t shatter.

I do not have a strong inclination either way: to say that the purported implicature of the consequence relation is or is not detached when (16) is substituted with one of (17)-(19). The reader may have stronger inclinations.

Fortunately for the critic of the conversational implicature version of the Material Implication Account, neither result is particularly helpful to such an account. If the purported implicature *does* detach, then this case fails to exhibit yet another feature of conversational implicature, since conversational implicature is non-detachable. However, even if the purported implicature *does not* detach, the failure to detach does not uniquely tell in favor of conversational implicature as opposed to being a part of the literal meaning of the utterance. I owe this point to Bach (1999), who points out that what is a part of the literal meaning of an utterance—what is said by the utterance strictly speaking—will of course not be detached when the meaning of the utterance is stated in equivalent terms. (See p. 335.) If the literal meaning of the utterance were not preserved, then the substituted terms would not be equivalent. So, the test of non-detachability is a faulty test when used to try to uniquely pick out instances of conversational implicature. The test does distinguish between conversational and conventional implicatures, since conventional implicatures are non-detachable (consider substituting ‘and’ for ‘but’), but it fails to distinguish between conversational implicature and literal meaning. Thus, the test for non-detachability does not support the account of the conditional as conversationally implicating the consequence relation any more than it supports an account according to which the consequence relation is part of the literal meaning of the conditional. Failing it this test would constitute yet another difference between the present case and cases of conversational implicature. Passing it would not tell in favor of conversational implicature over competing views that locate the consequence relation in the meaning of the conditional.

III. Conclusion

I have argued that any account of conditionals must take seriously the suggestion of a consequence relation between the antecedent and consequent of an indicative conditional. Since

such a relation does not exhibit the features of mere conversational implicature, when paired with the Material Implication Account, the conversational implicature version of the Material Implication Account is untenable. Such a position leaves open the possibility of a version of the Material Implication Account that either (i) does not attempt to explain the paradoxical nature of the paradoxes of material implication or (ii) attempts to explain it by means of conventional implicature. The first route is unfortunate, because the paradoxes of material implication are compelling and in need of explanation. The prospects for success of the second route cannot adequately be assessed here. However, I will mention that the very category of conventional implicature has met with some dubiousness,¹⁹ and it is unclear exactly why a locution with the conventional implicature of consequence would have the truth conditions of the material conditional. We should put to rest the idea that what explain the paradoxical nature of the paradoxes of material implication are pragmatic considerations arising from the assumption of cooperation with Gricean conversational maxims, and whether the Material Implication Account can be defended by other means must be left for discussion elsewhere.

¹⁹ See Bach (1999), p. 340 fn. 18, for an argument that what is purportedly conventionally implicated cannot be a part of the meaning of an utterance without participating in its truth conditions. DeRose (2002), pp. 197-198 fn. 16, also expresses wariness. See also Stanley (2002), pp. 12-13.

CHAPTER TWO

THE RELEVANCE TEST

Abstract: The Material Implication Account (MIA) of indicative conditionals holds that the indicative conditional ‘If P, Q’ in English has the same truth conditions as the material conditional ‘ $P \supset Q$ ’ in classical logic. Discussions of this account have tended to center on the so-called paradoxes of material implication and questions of assertability. I present a new challenge to the MIA, based on a test called the Relevance Test, whose judgment is independent of judgments about sufficiency for inference or assertability. I argue that the Relevance Test sets a very low standard that the MIA fails to meet.

The Material Implication Account of indicative conditionals holds that the indicative conditional ‘If P, Q’ in English has the same truth conditions as the material conditional ‘ $P \supset Q$ ’ in classical logic—that is, that it is false when P is true and Q is false, and it is true otherwise.¹ Discussions of this account have largely centered on the so-called paradoxes of material implication, which show that some inference involving ‘ \supset ’ is not sufficient for the corresponding inference involving ‘if’, and the numerous attempts to mitigate the paradoxes by introducing rules of assertion to explain their seemingly paradoxical nature. Such a discussion is helpful; and, should the Material Implication Account turn out to be the best theory, some story about the rules of assertion for ‘If P, Q’ must indeed be told. However, it may be that we need not settle the fate of the Material Implication Account in the court of assertability. In this chapter, I present an independent problem for the Material Implication Account. This problem is worth discussing not only because it is independently interesting but also because it does not turn on questions having to do with assertability or sufficiency for inferences. The discussion of the paradoxes and assertability, while

¹ Proponents of some version of the Material Implication Account include Russell (1960), Clark (1971), Dale (1974), Quine (1974), Lewis (1976), Jackson (1979), Grice (1989), and Barker (1997).

valuable, misses one of the most counterintuitive aspects of the Material Implication Account—namely, that it assigns truth values to conditionals for seemingly the wrong reasons. Even before judging the truth values assigned by this theory, we can see that its account of why a conditional has the truth value it has cannot be correct. In fact, the Material Implication Account assigns truth values to conditionals for reasons that seem to be *irrelevant* to the conditionals' truth values. Such a problem, since it is independent of considerations of both sufficiency for inference and assertability, is not open to pragmatic rejoinders of the same sort on behalf of the Material Implication Account.

In section I, I briefly discuss the two kinds of challenge to the Material Implication Account that have taken up the majority of the literature challenging that account: the paradoxes of material implication (which take two forms) and challenges to accounts of assertability conditions. (Readers familiar with these challenges may prefer to move directly to section II.) In section II, I introduce a test that any account of a linguistic construction in natural language must pass. I call this test the Relevance Test. Then, in section III, I argue that the Material Implication Account (hereafter, MIA) fails the Relevance Test. In section IV I address an objection based on a purported proof of the MIA. I conclude that the MIA, no matter how padded with Gricean implicatures and assertability conditions, suffers from a major flaw distinct from those most discussed in the current literature.

I. Two main kinds of objection to the Material Implication Account

The philosophical conversation about the Material Implication Account has centered on two main kinds of challenge to it.² The first challenge takes some valid inference involving ' \supset ' in classical logic and argues that the same inference is invalid for 'if...then' in English. Such inferences are

² Another challenge that has received a fair bit of attention has been the argument that indicative conditionals lack truth values. For presentations of this challenge, see, for example, Gibbard (1981), Appiah (1985), Edgington (1986, 1995, 2000), Bradley (2000) [See Douven (2007) for a rejoinder to Bradley.], and Bennett (2003), chapter 7. For arguments in favor of conditionals' having truth values, see, for example, Dale (1974), Lewis (1976), p. 84, and Chapter Four of this dissertation.

collectively referred to as the paradoxes of material implication. For example, ‘Q’ entails ‘ $P \supset Q$,’ but

(1) I will wake up before noon tomorrow.

does not entail

(2) If I die in my sleep tonight, I will wake up before noon tomorrow.

Numerous examples of problematic inferences can be generated. Discussion of the paradoxes of material implication takes place along two dimensions: sufficiency for inference and assertability. The problem as seen through the lens of sufficiency for inference is that there are formulae (which collectively let’s call ‘A’) whose truth is sufficient for some formula B in classical logic, where (i) ‘ \supset ’ appears at least once in A or in B and (ii) the truth of the English translation of A is not sufficient for the truth of the English translation of B.³ The paradoxes of material implication likewise appear in a closely related but distinct kind of challenge, having to do with assertability. Not only are there ‘ \supset ’-involving valid inferences in classical logic whose English-language counterparts are invalid; these English-language counterparts also fail to be highly assertable.⁴

In the face of the paradoxes of material implication, proponents of the MIA have often added a further component to the account, explaining in terms of rules for assertability the paradoxical nature of these inferences or their lack of assertability in English. Thus, the second main kind of challenge to the Material Implication Account consists in challenges to these rules for assertability.⁵ The details of these challenges vary, but what is important for our present purpose is that the discussions with respect to this kind of challenge all occur on the level of assertability rules and whether or not those rules can consistently be applied to yield intuitive results. The challenge I

³ For discussions of challenges of this kind, see, for example, Stevenson (1970), Clark (1971), Young (1972), Adams (1975), Gazdar (1979), Grice (1989), and Edgington (1995), pp. 244ff.

⁴ For a focus on the paradoxes in terms of assertability, see, for example, Jackson (1979) and Douven (2008).

⁵ See, for example, Jackson (1979), Strawson (1981), and Douven (2008).

present in this paper is not a challenge that the assertability rules are intended to address, and so discussion of the rules will not bear on the applicability of my results here.

II. The Relevance Test

I present a different kind of challenge to the MIA, making use of the Relevance Test:

The Relevance Test: An account of the truth conditions of a sentence should be such that only relevant factors determine its truth value.

I would expect the Relevance Test to enjoy a great deal of intuitive support, but here some things shall be said in its favor.

Employing the Relevance Test requires the assumption that, in general, even in the absence of an account of the truth conditions of a sentence, competent speakers are able to make good judgments about whether or not some fact is relevant to the truth value of that sentence. This assumption is a rather modest one—much more modest than requiring, for example, that competent speakers *know* the truth conditions of the sentences of the languages with respect to which they speak competently. There are some sentences in natural language whose truth conditions might be difficult for a competent speaker of the language to discern, even if the speaker understands the meanings of the words in the sentences and understands the ways in which these words are connected. Some sentences are too long or too convoluted or involve too much embedding of clauses and phrases for us to expect the competent speaker to be able to state their truth conditions. Furthermore, vague sentences are such that their truth conditions may be difficult to state. However, which facts are relevant to the truth of a sentence are, in general, something competent speakers should be able to discern even in the absence of an account of the truth conditions of that sentence. So, we who are competent speakers are licensed to make judgments about which facts are relevant prior to developing an account of the truth conditions for conditionals and then to use those judgments to evaluate various accounts.

It is worth noting that the non-conditional cases I mention above, in which the truth conditions may be difficult to know or to articulate, *do* seem apt for subjecting to the relevance test. In cases of vagueness, I may not know quite what it takes for

(3) David is bald.

to be true, but I know that hair facts are relevant. In the case of complicated sentences such as

(4) I don't know half of you half as well as I should like; and I like less than half of you half as well as you deserve.⁶

it may take some time for me to work out whether it is a compliment or not, but I at least know that facts about liking are relevant to its truth value. However, if the MIA is correct, the competent speaker is radically mistaken as to which facts are relevant to the truth values of indicative conditionals. So I argue in Section III.

One might argue that we should restrict the use of the Relevance Test to transparent contexts—that is, contexts in which one can substitute co-referring terms without changing the truth value of the whole—but this restriction is unnecessary, and showing why it is unnecessary serves to illuminate two important features of the Relevance Test. The first is that even tricky contexts such as opaque (that is, non-transparent) contexts are still clear enough for the Relevance Test to be applicable, because we can still tell what sorts of facts are relevant to the truth values of sentences involving opaque contexts. The second is that two differing accounts of the truth values of sentences can both pass the Relevance Test, showing how low a bar is set by the test.

It can be helpful to consider opaque contexts such as belief ascriptions in cases such as the Superman fiction, wherein Lois at some time does not know that her coworker Clark Kent is the same person as the superhero Superman. It can be difficult to discern the truth values of sentences such as (5):

⁶ Bilbo Baggins in Tolkien (1994), p. 29.

(5) Lois believes that Clark Kent can fly.

Is (5) true or is it false, given that Clark Kent is Superman and that Lois believes that Superman can fly? We might not know whether or not, within the fiction, (5) is true. But we do know that facts about Lois's beliefs are relevant to whether it's true: facts about how her beliefs relate to a certain person or to her other thoughts about that person, and what she thinks this person can do, as well as facts about how the parts of the sentence contribute to the whole. I submit that Lois herself would agree that these facts are relevant to the truth value of (5). Thus, even with sentences involving opaque contexts, it can be clear enough what facts are relevant to the truth value of the sentence for the Relevance Test to be applied. We need not restrict the test to sentences involving only transparent contexts. Of course, indicative conditionals do not in general involve opaque contexts, so my challenge to the Material Implication Account could proceed even if the Relevance Test were restricted to sentences involving only transparent contexts. I argue that the Relevance Test is useful even for opaque contexts, whose truth conditions are famously difficult to discern, merely to point out that we can expect to be able to discern what is relevant even in notoriously inscrutable cases. Thus, we can block the potential objection that perhaps we should not be expected to tell pre-theoretically what is relevant to the truth values of conditionals; we seem to be able to tell what is relevant even in these cases where truth values are known to be difficult to discern.

Furthermore, it seems that even differing accounts of (5) can both pass the Relevance Test. One account might say that (5) is true, because (5) is true just in case Lois believes of the person to whom 'Clark Kent' refers that that person can fly; and 'Clark Kent' refers to Superman, whom Lois believes can fly. Another account might say that (5) is false, because (5) is true just in case what 'Clark Kent' contributes to the sentence is something by which Lois attributes to Clark Kent the ability to fly; and 'Clark Kent' contributes to the sentence a sense such as *my coworker at The Daily Planet* by which Lois does not attribute to Clark the ability to fly. Both these accounts pass the

Relevance Test, because both use only features relevant to the truth value of the sentence to determine the truth value of the sentence—features such as Lois’s beliefs about Clark and what ‘Clark Kent’ contributes to the sentence (whether that contribution is a direct reference to the person himself or else a sense that represents that person under some guise). The bar set by the Relevance Test is so low that two differing accounts can both pass it. The failure of an account of the indicative conditional to pass the Relevance Test would thus count as a strong reason for rejecting that account, given how easy it is to pass, even by two conflicting accounts.

Before moving on to how it is that the MIA fares with respect to the Relevance Test, I want to address two potential concerns. First, it is important to note that failure of the Relevance Test really is distinct from being subject to the insufficiency challenge that I described in section I. It is well known that the MIA has met with challenges of insufficiency, but what guarantees that the failure of the Relevance Test is not just a symptom of insufficiency for inference? To see how my challenge is distinct from the insufficiency challenge, and in some ways more problematic, here is an example of an inference not involving conditionals that fails to be valid but which does not fail the Relevance Test. The fact that an account can fail according to the insufficiency challenge but not fail the Relevance Test shows that the two are distinct problems. Consider the following invalid inference from (6) to (7):

- (6) Hilary is a math major, and she has passed modal logic.
- (7) So, it’s false that either Hilary is a math major or Hilary has passed modal logic.

The invalid inference here is based on an account of the disjunction as exclusive, according to which ‘A or B’ implies ‘not both A and B’. Assuming that such an account is wrong, we have a failure of sufficiency for inference (or assertion) here. The truth of (6) does not guarantee the truth of (7). The inference from (6) to (7) fails to be valid, but the determination of the truth value of (7) based on this mistaken account does not constitute a failure of the Relevance Test. The fact that Hilary is a math major and has passed modal logic is relevant to whether or not the disjunction negated in (7) is

indeed false. In fact, the facts grounding (6) do determine the truth value of (7): namely, that (7) is false, contrary to the judgment of the account of the disjunction as exclusive. Thus, the challenge presented by failure of the Relevance Test is distinct from the failure to produce a valid inference.

Another potential worry is that, far from setting the low bar it seems to set, the Relevance Test asks too much of an account of a sentence's truth conditions. Perhaps there is a sense in which the truth conditions can be just those, and nothing more—not a guide to *why* a sentence is true or false, but just a guide to *whether* a sentence is true or false. Now, there are different kinds of explanation, and the Relevance Test tests only one. Suppose you walk into a kitchen where havoc has been wrought. There are at least two senses in which you may wonder why it is true that there is flour on the ceiling and egg on the floor. The linguistic explanation is that it is true that there is flour on the ceiling and it is true that there is egg on the floor, which suffices for the truth of the conjunction. (Such an explanation, of course, is not very helpful if what you want to know is how it came to be that those things were in those places at all.) This kind of explanation seems appropriate to expect of a sentence's truth conditions. Truth conditions tell us what determines the truth value. In some sense, they tell us what, linguistically speaking, accounts for the truth value of the whole. It is neither inappropriate nor burdensome to ask that an account use only relevant factors in determining the truth conditions of a sentence. In the case of the conjunction, no one denies that the truth value of 'There is flour on the ceiling' is relevant to the truth value of the conjunction. I think we will find that the analogue in terms of the MIA does not receive the same kind of intuitive support. It is to proving this result that I now turn.

III. How the MIA fails the Relevance Test

In this section, I argue that the MIA fails the Relevance Test. First, I need to establish the appropriateness of treating the Relevance Test as a test for an account of indicative conditionals,

given that many argue that indicative conditionals lack truth conditions. (See, for example, Adams (1975), Appiah (1985), and Edgington (1986).) The Relevance Test sets a standard for an account of the truth conditions of a sentence in natural language, and so it is appropriately used only on sentences that have truth values. If indicative conditionals lack truth values, then the test may be inappropriate. However, I am using the Relevance Test to argue against a particular account—the MIA—according to which indicative conditionals *do* have truth values. So, my challenge to the MIA relies on premises that the proponents of the MIA should accept, even if they are unacceptable to many of its detractors. Thus, the argument should be dialectically acceptable even to those who deny that indicative conditionals have truth values.

In order to establish that the Material Implication Account fails the Relevance Test, it will be helpful to deal with particular conditionals and then generalize from there. I present two cases in which (i) one does not know the truth value of a conditional, (ii) one is presented with new information, (iii) this new information seems irrelevant to one's pursuit of the truth value of the conditional, but (iv) it is information that the MIA treats as determining the truth value of the conditional. Consider (8):

(8) If Leslie is in Barcelona, then Leslie is in a former territory of the Roman Empire.⁷

As a competent speaker who does not know much about the former scope of the Roman Empire, I admit that I do not know whether (8) is true or false. Suppose I find out, while considering (8), that Leslie is in Pawnee, Indiana. The MIA, of course, deems (8) to be true in cases in which the antecedent is false, such as cases in which Leslie is in Pawnee. According to the MIA, the fact that Leslie is in Pawnee, Indiana, is relevant to the truth value of the claim that if Leslie is in Barcelona, then Leslie is in a former territory of the Roman Empire, because it is sufficient for the truth of the

⁷ Note that this conditional does not suffer from a relevance problem of a different kind. Its antecedent is relevant to its consequent.

conditional. But, intuitively, this information is not relevant. Because I lack the relevant knowledge of the Roman Empire, I do not know whether or not to endorse (8), even when it is pointed out to me that Leslie is not, in fact, in Barcelona. In order to know what to think about the truth value of (8), I want to know whether or not *Barcelona* is in a former territory of the Roman Empire. Leslie's actual location is not relevant at all! The conditional is true or false—though I do not know which—regardless of where Leslie is. In fact, I would think that the person who keeps pointing out to me Leslie's actual location was very mistaken indeed about what statements like (8) mean.

According to the MIA, the falsity of the antecedent is sufficient for the truth of the conditional. When we encounter a conditional with a false antecedent and false consequent (an FF conditional), we can ask, "Why, according to MIA, is this conditional true?" The simplest answer is that its antecedent is false. Every conditional with a false antecedent is true, according to the MIA. So, one does not need to consult the consequent of the FF conditional at all to determine its truth value. The reason the conditional is true is that its antecedent is false. Yet, in fact, what makes the antecedent false—and so, according to the MIA, what makes the conditional true—is irrelevant to the truth value of the conditional. Thus, the MIA fails the Relevance Test: it treats Leslie's actual location as relevant to the truth value of (8), when it is not relevant.

Many have pointed out that the mere falsity of the antecedent is not sufficient for the truth of the conditional in English. My argument goes beyond this claim to point out that, for conditionals with false antecedents, the fact that makes the antecedent false is in some cases not even relevant to the truth value of the conditional. Note that it is not more favorable for the MIA if we treat the dual falsity of the antecedent and the consequent, rather than the falsity of the antecedent alone, as that which determines the truth value of the conditional. For the fact that determines the falsity of both the antecedent and the consequent is the very same fact that we intuit to be irrelevant: the fact that Leslie is in Pawnee, Indiana. Enlarging our focus to whatever facts

make the antecedent false *and* the consequent false merely doubles the MIA's reliance on this irrelevant fact, overdetermining its failure of the Relevance Test. The foregoing example is merely one case that displays a more general problem with the Material Implication Account. The case of (8) relates to the paradox of material implication that points out that 'not-P' is not sufficient for (asserting) 'If P, then Q.' But the problem with failing the Relevance Test is not the same as a failure of insufficiency for inference or for assertion: the problem is that facts that are irrelevant to the truth value of (8) are treated as determining that truth value. Again, an account can suffer from failure of sufficiency for inference or assertion without failing the Relevance Test, so this challenge is distinct.

We can consider another example that displays the failure to pass the Relevance Test, where the closest paradox of material implication is of a different kind: 'Q' entailing 'If P, then Q.' Suppose that a scientist friend, Dikembe, shows you an object that is made of a newly-synthesized substance for which there seems to be some non-linear relationship between temperature and durability. At some temperatures it is much more durable than at others, but the relationship between temperature and durability has not yet been discovered. All that is known up to this point is that some increases in temperature lead to increases in durability, whereas some increases in temperature lead to decreases in durability, and no regular pattern has been established. It is Dikembe's work to establish such a pattern through testing the substance. Curious, and granting the MIA's assumption that conditionals have truth values, you wonder whether (9) is true or false.

(9) If Dikembe applies 50 psi to the substance at 100° F, then it will be crushed.

Dikembe asks you to make a guess about (9); and, not knowing any more information than is mentioned here, you guess that (9) is false. Now, suppose that Dikembe applies 100 psi of force on the substance at 30° F, and it is crushed. In that case, were you right or wrong about (9)? According to the MIA, you were wrong. The object has, in fact, been crushed, and so the consequent of (9) is true, rendering the whole conditional true according to the MIA. Has (9) been proven? No! For, as

you might rightly protest, this crushing of the object is irrelevant to the truth value of (9). (9) has been neither proven nor disproven, and the fact that the object has been crushed, which renders the consequent of (9) true, is irrelevant to whether (9) itself is true, because it has been crushed under a different amount of pressure at a different temperature, and there is possibly no relation between its being crushed in this circumstance and its being crushed in the circumstance described by the antecedent of (9). The new information about the crushing is irrelevant to the truth value of the conditional.

One might object here that the consequent of (9) contains a hidden indexical, thus: ‘If Dikembe applies 50 psi to the substance at 100° F, then *in that case* it will be crushed’. If such a hidden indexical is part of the meaning of the conditional, then the situation described is not a case in which the consequent is true, because *that case* has not obtained. Assuming the MIA, such a reading can be symbolized as follows: ‘ $P \supset (P \ \& \ Q)$ ’, for which there are two things of note. First, this new conditional is likewise true, according to the MIA, since the antecedent is false in the described scenario—and this represents a failure of the Relevance Test of the same nature as the failure in the case of (8). The fact that Dikembe has not yet tested the durability of the object at 50 psi and 100°F is not relevant to the truth value of the conditional. Second, ‘ $P \supset (P \ \& \ Q)$ ’ is logically equivalent to ‘ $P \supset Q$ ’. Thus, the fact—if it is one—that the consequent contains a hidden indexical does not affect the treatment of the conditional according to the MIA. The claim that you have made the wrong prediction could then be made in a multi-step process by the proponent of the MIA: (i) Q is true, so (ii) $P \supset Q$ is true; (iii) $P \supset Q$ entails $P \supset (P \ \& \ Q)$, so (iv) $P \supset (P \ \& \ Q)$ is true, so (v) you are wrong in claiming that (9) is false. In the cases of both (8) and (9), facts that are irrelevant to the truth value of the conditional are treated by the MIA as determining the truth value of the conditional. The MIA thus fails to meet the rather modest standard of the Relevance Test.

IV. Objection

In spite of the fact that the Relevance Test does not deal with sufficiency for inference, one might object to the rejection of the MIA on grounds related to sufficiency for inference. In particular, it seems that there is a simple proof of the Material Implication Account as follows: According to the MIA, 'If P, Q' is equivalent to 'not-P or Q'. It seems that 'If P, Q' can indeed be derived from 'not-P or Q'. First, assume P. By disjunctive syllogism with 'not-P or Q', derive Q. Then, by conditional introduction, you have 'If P, Q'. There is also an intuitive English-language version of the proof, though it is more natural in the (also-MIA-licensed) derivation of 'If not-P, then Q' from 'P or Q'. Suppose you know, of some crime, that either the butler did it or the gardener did it. It seems that, from that information alone, you can deduce that if the butler didn't do it, the gardener did it. If the disjunction entails the conditional, then the equivalence of the two is half-proven.

My response to this purported proof is not novel, but it is worth mentioning in light of my dismissal of the MIA. As others have noted, the disjunction seems like good evidence for the conditional only in cases in which the disjunction is not believed on the basis of believing one of the disjuncts to be true.⁸ If I believe that either the butler did it or the gardener did it on the sole basis of my belief that the butler did it, then I have no basis on which to infer that if the butler didn't do it, the gardener did it. If I become convinced that the butler didn't do it, I will deny the disjunction rather than accepting the consequent of the conditional. According to Jackson (1979), this phenomenon should be accounted for solely in my willingness or unwillingness to assert the conditional, rather than in my willingness or unwillingness to believe the conditional, but that claim seems wrong. It seems wrong to infer the conditional from the disjunction in cases in which I

⁸ Lycan (2001, p. 86), for example, makes an equivalent point with respect to inferring 'If A, then B' from 'not both A and not-B', which does not seem to be a good inference if one accepts the negated conjunction on the basis of disbelieving A.

believe the disjunction on the sole basis of believing one of its disjuncts. Despite the intuitiveness of the inference, upon reflection it is shown to be invalid in English.

V. Conclusion

Challenges to the Material Implication Account have centered on the paradoxes of material implication and the prospects for saving the account via a story about assertability. The Relevance Test presents a different challenge to the MIA in terms of what facts it deems relevant to a conditional's truth value. This problem is distinct from the question of whether inferences licensed by the MIA are made on sufficient grounds. A sentence P may be insufficient for Q while the mistaken truth conditions licensing this inference pass the Relevance Test, as we saw in the case of wrongly inferring 'not (A or B)' from 'A and B' based on mistakenly treating disjunctions as exclusive. Failing the Relevance Test is a problem that is distinct from problems with insufficiency for inference. The Relevance Test sets so low a bar that failing it is a serious problem. To paraphrase what someone once said about the Ten Commandments and human beings, the Relevance Test sets the bar on the ground, and the MIA trips over it.

CHAPTER THREE

TT CONDITIONALS

Abstract: Two phenomena present the need for a kind of flexibility in a successful account of indicative conditionals. First, there are some conditionals that have true antecedents and true consequents (TT conditionals) that are nonetheless false. Second, there are cases, such as successful conditional prediction, in which the mere truth of the antecedent and consequent *does* seem to render the conditional true. At least initially, these two phenomena appear to be in tension with each other, but a good account must accommodate both of them. In this paper, I argue that these phenomena count against various accounts of conditionals, including the Material Implication Account, various versions of the Possible Worlds Account, and various Suppositional/Probabilistic Accounts. I argue that the resources available to add flexibility to these views are unattractive.

Two kinds of phenomena present the need for a kind of flexibility in a successful account of indicative conditionals. First, there are some conditionals that have true antecedents and true consequents (hereafter, TT conditionals) that are nonetheless false or unacceptable—for example, ‘If grass is green, then Obama is President’. Few of the major accounts of indicative conditionals accurately predict this fact. A second phenomenon appears to be in tension with the first: in certain cases such as successful conditional predictions, the joint truth of the antecedent and consequent *does* seem to render the conditional true. (The tension between these two phenomena is merely apparent, for, as I will argue in Chapter Five, my own account accommodates both phenomena unproblematically.) If the analysis in this paper is correct, then just about every major account of indicative conditionals in the philosophical literature is mistaken with respect to one or the other of these two phenomena.

After discussing how most major accounts of indicative conditionals fail to allow for the possibility of false or unacceptable TT conditionals (section I), I argue that these two phenomena

exist (section II) and therefore count against various accounts of conditionals, including the Material Implication Account, various versions of Possible Worlds Account, and various Suppositional/Probabilistic Accounts. I argue that the resources available to add flexibility to certain accounts are unattractive (section III).

I. How the major accounts fail to allow for false or unacceptable TT conditionals

In this section, I show how intuitive versions of all three major accounts fail to accommodate the possibility of false or unacceptable TT conditionals. (In this paper I assume that indicative conditionals have truth values, except when specifically addressing accounts that claim the contrary, and so sometimes I will drop discussions of acceptability. However, as I shall argue, nothing hinges on this assumption. All my arguments can be made in terms of acceptability or assertability to accommodate accounts that deny that indicative conditionals have truth values. For a discussion of the relations among truth, acceptability, and assertability, see section I.C.) The Material Implication Account and Possible Worlds Account eliminate the possibility of TT conditionals altogether. The case with respect to Suppositional/Probabilistic Accounts is more complicated, but these accounts likewise fail to accommodate the kinds of conditionals that serve as support for the claim that some TT conditionals are false. The recourses available to proponents of these kinds of account (which some of their proponents take) are discussed in section III. The claim that all TT conditionals are true or acceptable is sometimes referred to as the Centering Principle, so called because of the practice in Possible Worlds Accounts of “centering” the possible worlds on the actual world in cases in which the antecedent is true.¹ Since this term fits best with Possible Worlds Accounts and less well with the other accounts, I prefer to frame the discussion in terms of TT conditionals.

¹ See Hajek and Hall (1994), p. 107, note 5.

A. The Material Implication Account

The Material Implication Account² most straightforwardly treats all TT conditionals as true. According to the Material Implication Account, ‘If A, then B’ has the same truth conditions as ‘ $A \supset B$ ’, which are the same truth conditions as ‘ $\sim A \vee B$ ’. Since, in a TT conditional, the consequent B is true, all TT conditionals are true according to the Material Implication Account.

B. Possible Worlds Accounts

The earliest and most intuitive versions of Possible Worlds Account treat all TT conditionals as true.³ According to a Possible Worlds Account, ‘If A, then B’ is true in world w just in case B is true in the closest (set of) A-world(s) to w , as determined by the appropriate selection function from w . Closeness of worlds is a matter of similarity. Since no world is more similar to the actual world than itself, cases in which A is true at the actual world will be cases in which the function selects the actual world as the closest A-world. Thus, if A and B are both true at the actual world, then ‘If A, then B’ is true at the actual world. Since A and B are arbitrarily chosen, this account has the result that all TT conditionals are true. There are variations of Possible Worlds Account that do not judge all TT conditionals to be true at the actual world because they operate on some set of propositions that is distinct from the actual world or because they select for multiple worlds even when the antecedent and consequent are both true at the actual world. I discuss these versions and whether or not these features help the accounts accommodate the relevant phenomena in section III.

² Proponents of the Material Implication Account include Grice (1989), Lewis (1976), and Jackson (1979).

³ Proponents of Possible Worlds Accounts include Stalnaker (1968, 1975, 1984, 2005), Ellis (1984), Pendlebury (1989), Weatherson (2001), and Nolan (2003).

C. Suppositional/Probabilistic Accounts

Suppositional/Probabilistic Accounts⁴ take seriously the Ramsey test⁵ for evaluating conditionals, according to which one should, in evaluating a conditional ‘If A, then B’, hypothetically add A to one’s stock of beliefs and then see whether or not B is also true. According to many Suppositional/Probabilistic Accounts, indicative conditionals lack truth values. While these accounts do not ascribe truth values to conditionals, and therefore cannot be guilty of ascribing truth to all TT conditionals, most of these accounts suffer from a structurally similar problem: they judge all conditionals whose consequents have a high subjective probability, given their antecedents, to be acceptable or assertable—which is false according to my analysis below. Suppositional/Probabilistic Accounts tend to deal not with the truth values of the antecedent and consequent, but rather with their subjective probabilities. Thus, when discussing Suppositional/Probabilistic Accounts, it will be useful to employ the notion of a BB conditional: a conditional whose antecedent and consequent are both believed, or, to put it differently, each of which has a high enough subjective probability to be something the speaker would be willing to assert in the relevant context. Since these accounts treat all conditionals with high conditional probabilities as acceptable or assertable, they treat most BB conditionals as acceptable or assertable. (The only BB conditionals that are not acceptable/assertable according to these accounts are those marginal cases in which the antecedent and consequent, though both believed by the speaker, nonetheless have a low conditional probability when combined. For example, I may believe that I will lose the lottery, and I may be highly confident that I have some small chance of winning, but ‘If I lose the lottery, I have some small chance of winning

⁴ Proponents include Belnap (1970), Adams (1975), Edgington (1995), Barnett (2006), and Douven (2008). Douven’s evidentialist account includes resources to avoid the consequence that all TT conditionals are acceptable, and I think those aspects of his account are successful. My complaint with Douven’s account is distinct: namely, its radically epistemic nature and its failure with respect to what Lycan (2001) calls Gallimore’s Problem (See Lycan (2001), pp. 69-72).

⁵ From Ramsey (1931).

it' has a very low conditional probability.⁶) This problem is structurally similar to the problem of judging all TT conditionals to be true for two reasons. First, TT conditionals that are believed to be TT conditionals are such that their component parts each have a high subjective probability. The reason such conditionals are treated as TT conditionals is that they are BB conditionals. So, as long as my examples of TT conditionals below are accepted as such, they also serve as BB conditionals. Second, there is a strong link between acceptability and assertability on the one hand and truth on the other. Truth entails acceptability and assertability of the kind I have in mind (defined below). The result of these two similarities is that the same examples can be used to show both the problem of treating all high-conditional-probability BB conditionals as acceptable/assertable and the problem of treating all TT conditionals as true.

I define acceptability as a quality of propositions or sentences such that they are deemed fit for private assent or rational endorsement. Assertability, on the other hand, is fitness for public assent or utterance. There are two things to note about the distinction between acceptability and assertability. First, a sentence can properly be categorized as acceptable or unacceptable even if it does not have a truth value. One may accept the sentence 'S' without accepting 'S' *as true*. For example, as Mellor (1993) uses the term, accepting a conditional consists in having a certain disposition. (See p. 238.) One may think of accepting a conditional, for example, as the disposition to come to believe the consequent upon coming to believe the antecedent, and this disposition does not require treating the conditional as true or false. Second, there are two levels of generality for both acceptability and assertability (the latter sometimes distinguished by spelling, as in Jackson (1987), p. 10)—acceptability/assertability in general and acceptability/assertability in a context—and only one of them is apt for use in the present discussion. Acceptability in general is fitness for rational endorsement in general, which most sentences will have just in case they are true (or, if they

⁶ Thanks to David Barnett for this example.

are not truth-valued, just in case they are objectively good sentences to endorse). Acceptability in general is a notion so close to truth that it may seem unnecessary; the only reason we need this notion is to account for the possibility that sentences such as conditionals could be fit to endorse, but not truth-valued. Acceptability in general is contrasted with acceptability in a context—that is, fitness for rational endorsement by a particular person in a particular epistemic state. Acceptability in general does not entail acceptability in a particular context, because the person in that context may not be in an epistemic state that is amenable to endorsement of the sentence in question, such as when the person lacks relevant evidence. For example, if I have good reasons to believe that my neighbor is at the grocery store, then the sentence ‘My neighbor is at the library’ is unacceptable to me, given my epistemic state, even if it happens to be true and so acceptable in general. Similarly, we can distinguish between assertability in general and assertability in a context. A sentence is assertable in general just in case it is acceptable. We can think of assertability in general as what would be assertable if conversational considerations such as relevance, brevity, etc., were not taken into account. A sentence’s assertability in a context, on the other hand, depends on contextual features such as its relevance and what pragmatic implicatures it would carry in that context. Assertability in a particular context can be diminished by features of the sentence that have little to do with whether or not the sentence is apt for rational endorsement, such as undue prolixity and bizarreness of content, even if the sentence is highly plausible and so acceptable. Thus, a great many sentences that are assertable in general are not assertable in some particular context. For example, ‘ $2 + 2 = 4$ ’ is assertable in general, but it is not assertable in very many contexts, because it is well-known and thus uninformative to most hearers. (Brian Weatherson (2014), p. 4, points out that there are problems with defining these notions such that assertability in a context entails assertability in general. It is possible for a sentence to be assertable in a context even if it is not acceptable to oneself [that is, assertable in general], for example in cases in which it is contextually or perhaps morally appropriate

to utter a lie.) I prefer to emphasize acceptability over assertability when discussing the analogue for truth in non-truth-conditional accounts, because it is easier to confuse assertability in general with assertability in a context. Unless I note otherwise, hereafter I will use ‘acceptability’ to mean acceptability in general and ‘assertability’ to mean assertability in general.

How is it that Suppositional/Probabilistic Accounts treat all conditionals whose consequents have a high subjective probability, given their antecedents, to be acceptable or assertable? These accounts have in common an endorsement of Adams’s Hypothesis. Adams’s (1975) Hypothesis takes the intuitive appeal of the Ramsey test and claims that the assertability of a conditional goes by the conditional probability of the consequent, given the antecedent. The general story is that, in a conditional ‘If A, then B’ whose antecedent and consequent both enjoy a very high subjective probability, A is already in one’s stock of beliefs, and so “adding” A to the stock does not change any of the thinker’s other beliefs, among which is the belief that B. Thus, B will also be believed under the supposition that A, rendering the conditional acceptable and assertable.⁷ However, this story needs some refining, because it is possible for A and B to each have high subjective probabilities, while the conditional probability of B, given A, is low, as in the lottery case mentioned above. Thus, I need to say more about how it is that Suppositional/Probabilistic Accounts treat all BB conditionals of the relevant kind as acceptable.

Whether or not Suppositional/Probabilistic Accounts treat all BB conditionals as acceptable depends on whether or not it is possible to have a scenario in which there is a BB conditional whose conditional probability—i.e., the subjective probability of the consequent, given the antecedent—is low. Intuitively, it seems as though it would be impossible to have a consistent assignment of probabilities such that the probability of each of A and B is high, while the conditional probability of

⁷ Others have noted that the treatment of all BB conditionals as acceptable is a consequence of Adams’s Hypothesis. See, e.g., Hajek and Hall (1994), p. 83.

B, given A, is low. It seems as though it would be inconsistent to believe both A and B to a high degree if the one entails that the other is *not* believable to a high degree. In reality, whether or not this scenario is possible depends on whether or not the probabilities of A and B are independent. In cases in which the probabilities of A and B are independent of each other, the probability of B, given A, is equal to the probability of B. (See Appendix One for the proof.) So if A and B have probabilities that are independent of each other and which are both high, it is impossible for their conditional probability to be low, and that means that all of those instances of BB conditionals come out as acceptable for Suppositional/Probabilistic Accounts.

However, it is possible for the probability of both A and B to be somewhat high while the probability of B, given A, is low—in cases in which the probabilities of A and B are not independent of each other. (See Appendix Two for an example.) In most standard conditionals that people have any occasion to use, the probabilities of A and B are not independent of each other, so why is it that it is only in marginal cases that BB conditionals end up being unacceptable according to Suppositional/Probabilistic Accounts? The reason is that, by definition, BB conditionals are conditionals in which the antecedent and consequent are not only considered to be highly probable by the person uttering the conditional; they are cases in which A and B are believed to be true. And cases in which (i) A makes B significantly less likely and (ii) you believe that both A and B are true are most likely cases in which you believe that A and B are true for independent reasons. Consider a case of buying a lottery ticket. Suppose you buy a lottery ticket with odds of 100 to 1 against winning, and you win. You know it is unlikely that you will win, given that you buy the ticket, but once you come to believe that you have won, the fact that a win was unlikely no longer gives you a reason not to believe that you won. You have other reasons to be confident in the fact that you won, even though your winning was unlikely. The prior probabilities of winning no longer inform your credence in the claim that you won. So it is with the BB conditionals I present in section II.A below.

It may be that something unlikely happens, but your credences are no longer sensitive to its unlikelihood once you believe that it has happened, such that you treat cases in which you believe both A and B as similar to cases in which the probability of the two is independent, and those are cases that the Suppositional/Probabilistic Account treats as acceptable. Just as I argue below that some TT (BB) conditionals are false, so it follows that they are unacceptable and are therefore unassertable. Thus, these accounts likewise fail to predict the first of our two phenomena below, since they give wrong judgments about the acceptability and assertability of the examples I give.

II. In need of flexibility

A. The first phenomenon

As I argued above, the major accounts of indicative conditionals each treat all TT conditionals as true or else all high-conditional-probability (that is, almost all) BB conditionals as acceptable or assertable. (In what follows, I sometimes gloss over the analogue with acceptability or assertability.) In fact, some philosophers have assumed that all accounts treat all TT conditionals as true. For example, Mackie (1973) states, without argument, that “non-material conditionals obey the top two lines of the standard truth table, though not the others” (p. 107). Now that the presumption in favor of all TT conditionals’ being true is established, I move on to discussing why this presumption is false.

This is the first of the two phenomena:

(P1) There are some false TT conditionals. (There are some unacceptable BB conditionals.)

In support of (P1), I discuss two kinds of false TT conditionals. Along the way, I present objections that motivate the introduction of new kinds of examples in support of (P1).

1. *False TT conditionals whose component clauses are irrelevant to each other.*

First, there are TT conditionals whose two component clauses are irrelevant to each other, such as (1)-(3).

- (1) If grass is green, then Obama is President.
- (2) If I exist, then the Seahawks won the 2014 Super Bowl.
- (3) If Obama is President in 2015, then the star Betelgeuse will someday go supernova.

Each of the above conditionals has a true antecedent and a true (or highly probable) consequent, and yet the above conditionals are intuitively defective. They sound wrong, misleading—perhaps even false. Although their antecedent and consequent clauses are independently highly plausible, putting these clauses together into the above conditionals produces sentences one might hesitate to assent to. Compare conditionals (1)-(3) with the following 'TT' conditionals, whose component clauses likewise enjoy a great deal of plausibility:

- (1') If grass is green, then grass is the same color as algae.
- (2') If I exist, then something exists.
- (3') If Obama is President in 2015, then healthcare will be a major subject of public discourse in the U.S. that year.

In contrast with (1)-(3), the conditionals (1')-(3') are plausible and easy to assent to. For those who ascribe truth values to indicative conditionals, (1')-(3') seem true, and the defective-soundingness of (1)-(3) stands out in stark contrast. In fact, conditionals (1)-(3) sound defective *as conditionals*. In each of (1)-(3), the consequent does not seem to be conditional on the antecedent at all. In this way, conditionals (1)-(3) are similar to speech act conditionals such as Austin's⁸ biscuit conditionals ('There are biscuits on the sideboard, if you want some.') and to semifactuals, which are conditionals admitting of or containing 'even if' or 'still' ('Even if I am feeling ill, I will not miss the ceremony.'). In the case of the biscuit conditional, the placement of the biscuits is not conditional on your desire.

⁸ See Austin (1961).

In the case of the semifactual, my attendance is not conditional on my feeling ill. In fact, the point of the semifactual seems to be to deny that the events described in the antecedent will have the expected effect on the truth value of the consequent. The case is similar with conditionals (1)-(3): the consequent is not conditional on the antecedent.

One might object here that it is unfair to use the concept of one thing's being conditional on another in a natural way and then apply it to sentences that happen to be (perhaps mistakenly) classified as conditionals, and there may be something to that objection. Nonetheless, we are discussing an account of standard conditionals, of the core of what we are up to when using 'if'-sentences, and it would be very surprising indeed if the true account of these sentences did not involve one thing's being conditional on another in this natural sense. In any case, it seems as though the consequent's not being conditional on the antecedent should count against the acceptability of these conditionals to some extent, even if this consideration is not decisive.

Note that, while bearing some similarity to semifactuals, conditionals (1)-(3) are not semifactuals. Adding 'even' before 'if' (and 'still' where appropriate) does not create a felicitous-sounding conditional. Following Levinson (1983),⁹ I will make note of an anomalous sentence by preceding the name of the sentence with a question mark.

?(1'') Even if grass is green, Obama is President.

?(2'') Even if I exist, the Seahawks still won the 2014 Super Bowl.

?(3'') Even if Obama is President in 2015, the star Betelgeuse will someday go supernova.

The locution 'even if' is used to express something that is contrary to expectations. It is appropriate in cases such as the following, wherein one would treat the truth of the antecedent to be a good reason for expecting the consequent to be false or as a possible means of making the consequent false:

(4) Even if am ill, I will still attend the ceremony.

⁹ See p. 6, fn. 3.

(5) Even if you cut the cord there, the 'TNT' will still detonate.

The conditionals (1)-(3), on the other hand, do not thwart expectations. There is no expectation that the greenness of grass is a good reason for believing Obama not to be President. Thus, these conditionals are not semifactuals. One might respond that the denial of expectations is merely pragmatically implicated by the addition of 'even' and 'still' to semifactuals, and so the fact that these conditionals do not involve denied expectations may not definitively exclude them from being semifactuals. My response is this: I think we have good reasons for thinking that how 'even' and 'still' work in semifactuals is through conventional implicature, not mere pragmatic or conversational implicature. As in the coordinating conjunction 'but', the connotation of thwarted expectations cannot be canceled, but this cancelation would be possible if the suggestion of thwarted expectations were merely a result of conversational implicature. Furthermore, even if doubt is cast on the diagnosis of semifactuals as necessarily involving thwarted expectations, the oddness of adding 'even'/'still' to these conditionals, as expressed in (1'')-(3''), is enough to show that (1)-(3) are not semifactuals. It is important to note that these conditionals are not semifactuals, because there are good reasons for treating semifactuals as non-standard and thus as poor rulers against which to measure an account of conditionals. If these conditionals were semifactuals, it might be appropriate to give them a slightly different treatment from standard conditionals and explain this different treatment in terms of how this kind of locution deviates from standard conditionals. (See Introduction, section I.B.) Since they are not semifactuals, (1)-(3) deserve no such special treatment on this basis.

(1)-(3) cannot be excluded from use as test cases for an account of conditionals on principle of being instances of some non-standard kind of conditional, and so an account of conditionals must accommodate the phenomena surrounding them within its treatment of standard conditionals. Obama's being President is not conditional on grass's being green, and so (1) comes off sounding

wrong, defective, as do (2) and (3) for similar reasons. I submit that the reason conditionals (1)-(3) sound wrong is that they are false. It is not the case that, *if* grass is green, then Obama is President. In what follows, I will write as though these 'TT' conditionals are false, but the phenomenon of (P1) still presents a problem even if it is merely the case that these conditionals are unacceptable and so unassertable. As I argued above, the major accounts of conditionals all either judge every 'TT' conditional to be true or else judge every high-conditional-probability BB conditional to be acceptable/assertable. Since (1)-(3) are at least unacceptable/unassertable, if not false, and since each of them also has a high conditional probability, they present a problem for every major account of indicative conditionals.

There is a reason beyond their intuitive unacceptability and non-conditional content for thinking that conditionals such as (1)-(3) are unacceptable (and so, false, if they have truth values). I believe the kernels of the following reasoning come from Lycan (2001), though the potential response and further reasoning are mine. Lycan points out that conditionals whose antecedents and consequents are irrelevant to each other are not only misleading to assert but also misleading to accept or believe privately. The reason is as follows: conditionals, whatever their truth conditions or literal meaning, at least seem to suggest some sort of connection between the antecedent and the consequent. We often use conditionals to suggest such a connection at least pragmatically when we assert conditionals, even if it is not literally a part of the meaning, and conditionals carry the same connotation in private reasoning. Suppose I accept a conditional on the mere basis of believing that its antecedent and consequent are both highly probable. Later on, I might retain acceptance of this conditional without remembering why I came to have that belief. If I have forgotten why I came to believe the conditional, then I might be liable to reason on the basis of that conditional in ways that are unsupported by the conditional, because the conditional was not believed on the basis of some apparent connection between the antecedent and the consequent. For example, I might treat the

conditional as apt for modus tollens reasoning, should I come to believe its consequent to be false, when a better response upon believing the falsity of the consequent would be to give up belief in the conditional itself. Consider the example of (3), ‘If Obama is President in 2015, then the star Betelgeuse will someday go supernova’. Suppose I come to believe (3) on the mere basis of believing both its antecedent and consequent to be true or highly probable. Decades later, I retain belief in the conditional while forgetting why it is that I came to believe the conditional in the first place. I then learn, due to some recalculation on behalf of astronomers or some unexpected cosmic event, that it is very unlikely that the star Betelgeuse will ever go supernova. Given my acceptance of (3) and modus tollens, I might be tempted then to conclude that it is very unlikely that Obama was President in 2015. However, since the reason for my initial acceptance of (3) was, in part, belief in the truth of its consequent, finding out that its consequent is false should lead me to reject (3) rather than to disbelieve its antecedent. Yet this reasoning is no longer available to me, and so I have been misled by accepting (3).

As I see it, two intuitive lines of response are available to someone who thinks all TT conditionals are true (or all high-conditional-probability BB conditionals are acceptable). First, one might respond that for practical reasons people should not expect there to be a connection between the antecedent and consequent of the conditionals they have accepted. It may seem that, in the case above, a rational person who finds out that the consequent is false should consider independently whether or not it is likely that there is some connection between the antecedent and the consequent such that modus tollens or some other reasoning is appropriate. If an appropriate connection is likely, then the reasoning may proceed. Yet this response undermines much of the value of accepting conditionals for use in further reasoning. If the conditional itself should not be used in further reasoning—but rather the thinker should look for some other fact to reason from instead, such as the truth values of the component clauses or else some fact about a connection between the

antecedent and the consequent—then accepted conditionals are rendered rationally inert. Accepting or rejecting conditionals is of no consequence, if one must always substitute the conditional with some other belief while reasoning.

Second, a doubter of (P1) may suggest, on the other hand, that one should not accept conditionals merely on the basis of believing their antecedents and consequents to be true, even though according to this person there is nothing defective or unacceptable about such a conditional. Such a suggestion may be germane practically speaking, but it serves to highlight how far such a treatment of conditionals is from how real people *use* conditionals. This suggestion has as a result that the primary use of conditionals—to express connections between events, ideas, etc.—is a more specialized use that, if it cannot be avoided, necessitates privately rejecting many conditionals that are no less true, meaningful, or acceptable than the ones that *can* be recommended for private acceptance. Furthermore, such a treatment of conditionals regards conditionals such as (1)-(3) as only nominally acceptable. For, according to this response, conditionals for which there is not a connection between the antecedent and consequent are not really fit for rational endorsement by most people. These conditionals are misleading, and so, practically speaking, they should not be accepted. Yet the whole point of the response was to defend these conditionals as acceptable in light of the charge that they are misleading. This line of response fails to defend them from that charge by judging them to be unacceptable after all.

So far, the discussion of (P1) has centered on the first kind of TT conditional, as exemplified by (1)-(3). The second kind of false TT conditional is motivated by a potential objection to my use of conditionals whose component clauses are irrelevant to one another. One might deny (P1) by arguing that the fact that examples (1)-(3) seem defective or false stems from the fact that they are merely odd—and not from falsity.¹⁰ After all, one has little reason to think that the greenness of

¹⁰ Thanks to David Barnett for pressing me on this point.

grass has anything to do with Obama's presidency, and so examples such as (1) come off as bizarre, and similarly for (2) and (3). If the fact that examples (1)-(3) sound defective can be explained by their oddness rather than their falsity, then they no longer support (P1), and if (P1) is unsupported then I have not successfully argued that accounts of indicative conditionals must account for (P1). This sort of strategy can be employed in defense of any of the major accounts, since it does not rely on details from any particular account.

I take this response to be an intuitive one, though I think that it does not succeed. In my own case at least, it is not merely a knee-jerk reaction to the oddness of (1)-(3) that underlies my judgment of them as false. Upon careful reflection, including awareness of their odd content, it still seems that in each of (1)-(3) the consequent is not conditional on the antecedent, and this fact seems to be a good reason for rejecting (1)-(3). These cases seem defective *as conditionals*, not merely defective insofar as they have odd combinations of subject matter. Furthermore, we can test whether or not the oddness of the subject matter in (1)-(3) is interfering with judgments about truth or falsity by considering whether or not the component clauses have the same effect when they appear in other linguistic arrangements. For example,

- (1a) Grass is green, and Obama is President.
- (1b) Grass is green, or Obama is President.

(1a) and (1b) are indeed odd, due to the combination of mutually irrelevant subjects, but upon reflection they can be heartily assented to without reservation. The oddness of the subject matter remains, but in the cases of (1a) and (1b), it is not accompanied by a hesitance to endorse the resultant sentences. It is the conditional alone that is difficult to assent to, which indicates that (1), for example, is defective in a way that is not wholly due to its odd content. Such a test shows that the fact that the two component clauses are irrelevant to one another does not affect one's judgments about the truth values of the whole in cases of conjunction and disjunction. If the oddness of the subject matter is enough to make the conditional seem difficult to assent to, when it

does not do the same for the conjunction or disjunction, there must be some aspect of the meaning of the conditional that renders the irrelevance inappropriate, not just in a particular context, but in general. Odd content is not enough to render a sentence unacceptable in general, and so this test presents evidence that the meaning of a conditional requires some connection between the antecedent and consequent. Thus, oddness of subject matter is not an excuse to dismiss (1)-(3) as examples of false TT conditionals. Rather, it gives us an indication that there is something distinctive about indicative conditionals, as compared with other kinds of sentences.

2. False TT conditionals that do not include irrelevance

Nonetheless, for those who are not convinced, I present a second kind of false TT conditional. These are certain kinds of case, bearing some artificial similarity with Gettier cases,¹¹ in which reality is contrary to expectations. Consider the following conditional:

(6) If it is windy, then I will hear the wind chimes chiming.

There is nothing odd or defective-sounding about (6) on its surface, and its antecedent and consequent are relevant to each other. In fact, (6) may have a great deal of initial plausibility. Nonetheless, there are possible scenarios in which (6) is a TT conditional and yet false. For suppose that it is windy, and that I can hear the wind chimes chiming, such that (6)'s antecedent and consequent are both true. Now comes the twist: the reason I can hear the wind chimes chiming is that someone is cleaning the wind chimes in the kitchen, where the wind does not blow. Such a scenario does not render (6) odd-sounding, but it does seem to be a scenario in which (6) is false, despite the fact that its antecedent and consequent are both true. In the envisaged scenario, the details of what makes the consequent true undermine the plausibility of the conditional. Someone

¹¹ Gettier cases are cases in which a person has a justified true belief, but the belief is true for reasons one would not expect. My cases are cases in which a conditional has a true antecedent and true consequent, but one of them is true for reasons one would not expect. See Gettier (1963).

who knows why I can hear the wind chimes chiming will hesitate to assent to (6), even if that person also knows that it happens to be windy, because in the envisaged scenario it is saliently plausible that I might not hear the wind chimes chiming while it is windy—for example, after the person who brought them inside sets them down on the kitchen counter. (6) seems wrong in this scenario, while not seeming odd in the same way as (1)-(3), and so it supports (P1) without the use of conditionals whose component clauses are irrelevant to one another.

Here I will discuss and respond to two possible objections to my reasoning about (6) that would serve to undermine my use of it as making trouble for the major accounts. First, one may read examples such as (6) and the scenario described above and not feel that (6) is best judged to be false (or unacceptable) in this scenario, but rather one may feel a kind of deep ambivalence about (6). It may not seem right either to heartily assent or else strongly dissent from (6) in the envisaged scenario. In this case, it may seem that conditionals such as (6) are borderline cases that should not be used as test cases for a theory. A good theory will explain all the central cases of some phenomenon and give plausible judgments for cases where we lack a strong pre-theoretical judgment. Perhaps (6) is just such a borderline case. If this is so, then it should not count against the major accounts that they give a judgment on (6) that was not pre-theoretically expected.

I take it that a reaction of deep ambivalence to cases such as (6) is understandable and may be common. I can sometimes get myself into a frame of mind in which this ambivalence seems more appropriate than does judging (6) to be false. Yet if this reaction is understandable and common, then a good account of indicative conditionals should explain one's ambivalence about cases such as this one. I think proponents of major accounts would be hard-pressed to explain this ambivalence by means of the features of their accounts. There is nothing particularly odd or unassertable about (6) on its face, and as I noted it enjoys a great deal of initial plausibility. Yet the ambivalence cannot adequately be explained as an instance of expecting one outcome and receiving

another; for in cases where a sentence has a high subjective probability and then one's initial probability judgment is challenged, there is in general no ambivalence about the result. I discuss the ability of my own account to explain reactions of ambivalence in Chapter Five, and if that discussion is successful then my account has an advantage in one respect, even if it turns out that a common response to (6) is ambivalence.

A second possible objection to my reasoning about (6) is to claim not that (6) is false or that one is appropriately ambivalent about (6), but that (6) is nonetheless true in the envisaged scenario due to the fact that the antecedent of the conditional makes the consequent likely. Perhaps the above reasoning treats (6) with undue strictness, expecting the conditional to be true only if the antecedent *necessitates* the consequent. Yet if someone asks you, 'If it is windy, will you hear the wind chimes chiming?', an answer of 'Probably' seems appropriate while also more or less amounting to endorsing the conditional. If it is windy, I will probably hear the wind chimes chiming—and that is true even though sometimes the wind chimes are taken down to be cleaned inside. Someone leveling this objection might complain that this case misleads the reader by focusing on a scenario in which one's reason for endorsing the conditional is undermined, when really the subjective probability of the consequent, given the antecedent, is still high enough to make the conditional acceptable.

Responding to this second objection requires producing the following sort of case: a false or unacceptable 'TT' conditional where the antecedent does not make the consequent highly probable, even though the two are relevant to one another. In such a case, we have a 'TT' conditional that is false—thereby supporting (P1)—while not being subject to the complaint either that its component clauses are irrelevant to one another or that it is true after all due to its antecedent's making its consequent probable. Cases of misconceptions about causal connections or probabilities can provide good example cases. Consider conditional (7):

(7) If Leslie is bitten by a tick, she will contract Lyme disease.

In our envisaged scenario, let us make (7) a TT conditional. Leslie is bitten by a tick, and she contracts Lyme disease. Ticks are known to carry Lyme disease. However, in this scenario let Leslie contract Lyme disease by some means other than the tick bite—for example, because she receives a blood transfusion from someone who had Lyme disease.¹² Here again we have a TT conditional, and its antecedent and consequent are relevant to each other, given the fact that tick bites transmit Lyme disease and are in fact the only known way this disease has been transmitted in the past. Another fact is important for making this case relevantly different from the case of (6): very few tick bites transmit Lyme disease. According to Paul Mead, who is the chief of epidemiology and surveillance activity at the CDC, “In areas where [Lyme disease] is very common, one out of every four or five ticks might be infected. In other areas where it’s much rarer, that may be more like one in 100” (Klein (2013)). This fact makes the case of (7) immune to the second objection that was leveled at (6): the charge that, despite the facts undermining the reason for endorsing the conditional, nonetheless it is acceptable due to the high probability that the consequent will obtain, given the antecedent. (7) is not such a case; the overwhelming majority of ticks do not carry Lyme disease, and so being bitten by a tick does not render it highly probable—nor even more probable than not—that Leslie will contract Lyme disease.

One might object that what matters in this case is not rendering the probability that Leslie contracts Lyme disease high, but rather *raising the probability* that Leslie will contract Lyme disease at all. Being bitten by a tick does raise the probability that Leslie will contract Lyme disease, and so according to the present suggestion, (7) comes out as true, in which case it is not a false TT conditional and does not support (P1). I see two problems with challenging (P1) by means of that approach. First, that approach is not likely to go along with rejecting (P1), since there will be many

¹² According to the Center for Disease Control, contracting Lyme disease in this way is medically possible, although there are no known cases in which contraction of Lyme disease by this means has occurred. See <http://www.cdc.gov/lyme/transmission/index.html>.

instances of TT conditionals (for example, (1)-(3)) where the truth of the antecedent does not raise the probability of the truth of the consequent, and so these TT conditionals turn out false according to this suggestion, vindicating (P1). Second, there are independent reasons for thinking such a condition cannot be necessary for the truth of a conditional. Consider, for example, true conditionals whose antecedents and consequents are both necessarily true, such as (8):

(8) If four is greater than three, then four is greater than two.

The conditional (8) is true, because any number greater than three is also greater than two. However, the consequent of (8) is necessarily true, and thus the subjective probability of the consequent independently of the antecedent is 1, which cannot be raised. Since the truth of its antecedent does not raise the probability of the truth of its consequent, (8) comes out as false according to the present suggestion. Since, as I think, (8) is true, this case shows that the present suggestion cannot be a necessary condition on a true conditional.

At this point proponents of the second objection may very well insist that even in this case, despite what one might be inclined to think, (7) is still true or, upon reflection, acceptable—as I shall insist the contrary. If this clash of judgments is what discussion comes to, then we may, to a small extent, be satisfied. For if I have done nothing but show that there are responsible, understandable readings according to which some TT conditionals are false (even if they are not universally shared), then I have shown that the analysis in this paper is worth considering along with the accounts that fail to accommodate it.

Despite the overwhelming support for treating all TT conditionals as true, I am not the first to have noticed that such a practice is problematic. For example, Pendlebury (1989), in discussion of the Ramsey Test, says “no ordinary speaker of English would accept any old belief which is independent of P as part of a justification for the claim that $P \rightarrow Q$ is true just because this belief is a reason for accepting Q” (203). In other words, the judgment that all TT conditionals are true does

not fit well with what ordinary speakers consider to be good reasons for believing a conditional, since none would believe a conditional merely based on the truth of the consequent. Ellis (1978) expresses similar doubts that all TT conditionals are true, although he falls short of claiming, as I do, that some are false. His complaint on p. 109 is that a principle that treats all TT conditionals as true “forces us to count such sentences as ‘If the weather is fine today, then I had bacon and eggs for breakfast yesterday’ as true, if both antecedent and consequent are accepted as true, which is, perhaps, counterintuitive.” Despite its counterintuitiveness, the bare bones of the three major accounts of indicative conditionals all judge every TT conditional to be true (or make a similar mistake with respect to BB conditionals and acceptability/assertability), contrary to (P1).

3. *A note on ‘whether or not’*

Before moving on to (P2), I would like to address one potential objection to (P1). The objection is that treating some TT conditionals as false fits ill with the seeming equivalence between ‘if’ and ‘whether’. More particularly, it might seem that ‘B whether or not A’ entails ‘If A, then B, and if not-A, then B’, which entails ‘If A, then B’. Yet if some TT conditionals are false, this entailment does not hold. For example, according to my analysis, ‘Obama is President whether or not grass is green’ does not entail ‘If grass is green, then Obama is President’, because the former is true while the latter is false. The purported entailment from the ‘whether or not’ sentence to the ‘if’ sentence seems intuitive, and so an explanation must be given for why it does not hold.

One reason for thinking that ‘B whether or not A’ does not entail the standard conditional ‘If A, then B’ is that, in many cases, ‘whether or not’ sentences seem equivalent to the conjunction of two conditionals, one of which is standard and the other of which is a semifactual. For example, ‘I will attend the ceremony whether or not I’m sick’ seems to have the force of ‘If I’m not sick, I will attend the ceremony, and even if I am sick, I will still attend the ceremony’. The second conjunct

there is a semifactual. The ‘whether or not’ sentence has the same force as the conjunction with the semifactual, because the ‘whether or not’ sentence seems intended to communicate that being sick will not have the expected consequence on my attendance, and the semifactual communicates this same message. Since semifactuals are not standard conditionals and do not entail standard conditionals, the fact that the ‘whether or not’ sentence seems best expressed by the conjunction of a standard conditional and a semifactual might serve as an indication that ‘B whether or not A’ does not in general entail the standard conditional ‘If A, then B’.

The cases of (1)-(3), whose component clauses are irrelevant to each other, do not generate examples of this kind. There is no presumption that the color of grass would have any bearing on whether or not Obama is President, and so the related ‘whether or not’ sentence seems merely to express that the fact that grass is green is irrelevant to the fact that Obama is President. However, all ‘B whether or not A’ cases are cases of asserting B, which then serves as the consequent of the purportedly-entailed conditionals. As the above reasoning shows, asserting ‘If A, then B’ on the basis of believing B is problematic. It is potentially misleading, even to oneself, to accept ‘If A, then B’ merely on the basis of believing B, because accepting the conditional for that reason opens one up to a problematic modus tollens in the case in which one remembers the conditional while forgetting the reason for believing it. Believing ‘If A, then B’ on the basis of B also seems to lead to accepting conditionals that are simply unacceptable, such as the ‘TT’ examples above or this example due to Dunn and Restall (2002): ‘If one scares a pregnant guinea pig, then all of her babies will be born tailless’.

Since speakers may not recognize the difference between a standard conditional and a semifactual, nor do they have incentives to be particularly strict about their reasons for accepting conditionals, it is understandable that ‘B whether or not A’ might be thought to entail ‘If A, then B’. Yet there are at least a couple reasons, in addition to all the reasons for thinking some ‘TT’

conditionals are false, for thinking that this purported entailment does not hold. The result is that what seemed to be omniconditional—to be true no matter what the conditions—ends up being unconditional. If it is true that ‘B whether or not A’ does not entail ‘If A, then B’, then conditionals may not be as logically weak as they have sometimes been thought to be. Perhaps because of the early influence of the Material Implication Account or the widespread treatment of all TT conditionals as true/acceptable, it has long been thought that ‘A and B’ and ‘B whether or not A’ both entail ‘If A, then B’. Yet if the analysis in this chapter is correct, the conditionals say both less and more than the other locutions: less because they do not assert their component parts, and more because what they do assert is not entailed by those locutions. The details of what conditionals assert will have to wait for another chapter.

B. The second phenomenon

The phenomenon of (P1) is enough to cast serious doubt on the major accounts of conditionals. But before I address more of the possible responses on behalf of the major accounts, I would like to point out another phenomenon surrounding TT conditionals, (P2), which further constrains a good account of indicative conditionals.

(P2) There are some cases in which the fact that the antecedent and consequent come out true *does* seem to make the whole conditional true (or acceptable), even when the consequent’s truth is not entailed or necessitated by the antecedent.

Our second phenomenon is exemplified in cases such as successful conditional predictions.

Consider (9).

(9) If Sophia rolls the die, the die will come up six.

If the events described in the antecedent and consequent both obtain, then the utterer of (9) is vindicated. She has made a true prediction. Assuming that the die is fair, the roll could turn out any of six different ways, but it happens to turn out in the one predicted by (9). This fact seems to be

enough to make (9) true. If, on the other hand, the antecedent and consequent are not both true, then (9) is false. It is the truth of the antecedent and the consequent that seems to account for the truth of the whole conditional. Such analysis generalizes to any case of straightforward successful conditional prediction: the conditional prediction seems to be true just in case—and, plausibly, *because*—the antecedent and consequent both come out true.¹³

It is important to note the difference in intuitive responses to (9), the case of clear successful prediction, and to (6) and (7), the cases that seem not to be straightforwardly successful, and are plausibly not successful at all. While responses to these cases will no doubt vary from person to person, or from time to time as one considers them in different lights, even a modest difference in how one responds to these kinds of case—(6) and (7) on the one hand and (9) on the other—is enough to generate the need for an explanation of the difference. One need not think, as I do, that (6) and (7) are false while (9) is true to motivate a consideration of the difference between these kinds of case. Something as mild as an unhesitating assent to (9)'s being true or acceptable or successful, combined with some ambivalence or hesitance in the cases of (6) and (7), still presents the need for explanation as to why these cases differ in acceptability. An account of indicative conditionals that anticipates and explains these differences will therefore have an advantage in that respect over accounts that do not accommodate these phenomena.

In support of the claim that the account of conditionals should do the work of explaining the difference between these two kinds of case, I want to give reasons for thinking that a certain

¹³ Lycan (2001) denies that successful conditional predictions such as (9) are true, though he thinks the corresponding backward-looking conditional, once the outcome is known to all, is true. Lycan's reason is that, before the die is rolled (or coin is flipped, in his example), alternatives are still "real possibilities" (that is, saliently epistemically possible). This response relates to Lycan's hesitance to make his theory such that reality, as opposed to the utterer's beliefs, constrains the interpretation of the conditional. According to his account, even if it is physically impossible for x to occur, if x is a "real possibility" for the utterer of a conditional, then x constrains what that conditional means and its truth conditions. Likewise, if y is not a "real possibility" for the utterer, no law of nature such that y must occur will affect the meaning or truth value of the conditional. As a result, Lycan's account suffers from Gallimore's Problem. (See Chapter Four for a discussion of Gallimore's Problem.)

non-conditional feature of the accounts cannot serve as an explanation. It might seem natural to suppose that the difference between the two kinds of response to these sets of cases consists merely in the fact that a false belief could serve as the reason for endorsing (6) and (7)—for example, the false belief that the wind chimes are outside for (6) and the false belief that most tick bites transmit Lyme disease for (7)—whereas none serves to underlie (9). This response says that it is our disapproval of this false belief that casts doubt upon the conditional in the first two cases. I think this response is mistaken. For we can easily add to the case of (9) above that the person makes the prediction because of the belief that there is a ghost in the room who wants the die to come up six. Let it be the case that there is no such ghost. The prediction seems successful nonetheless, and the person who bets on (9) has still straightforwardly won the bet.

Cases such as (9) are important not only because they present an interesting tension when combined with (P1) but also because they rule out a strict account of indicative conditionals. A strict account of indicative conditionals holds that ‘If A, then B’ is true just in case A necessitates B. Proponents of strict accounts include Mayo (1957). (See pp. 301-302.) The cases discussed in support of (P1) may seem to support a strict account of indicative conditionals. It may seem to be the fact that there’s some gap between the events described in the antecedent’s obtaining and the events described in the consequent’s obtaining that accounts for their unacceptability. But cases such as (9), in support of (P2), show that the antecedent’s necessitating the consequent is not required for a true conditional.

(P1) and (P2) seem to be in tension with each other. (P2) is exemplified by cases in which the truth of the antecedent and consequent seems to determine the truth of the conditional, and yet (P1) tells us that the truth of the antecedent and consequent do not suffice for the truth of an indicative conditional. Why is the truth of the antecedent and consequent sufficient in some cases but not in others? I submit that a successful account of indicative conditionals must accommodate

both (P1) and (P2). What we need is an account with flexibility—one that is not so loose as to accept all TT conditionals as true (since some appear to be false), but not so strict as to require that the antecedent necessitate the consequent (since sometimes the truth of the antecedent and consequent does render the conditional true in cases in which the antecedent does not necessitate the consequent).

III. How some accounts try to accommodate (P1) and (P2)

In this section, I discuss three resources available to different accounts that would help them to accommodate both (P1) and (P2), showing these resources to be inadequate.

A. Take refuge in the “Assert the Stronger” principle.

Proponents of the Material Implication Account such as Grice shift the discussion to talk of assertability, taking refuge in the “Assert the Stronger” principle. (See Grice (1989), pp. 61-62.) According to this principle, one should be maximally informative, though not to the point of violating rules of brevity. Someone who believes that ‘If P, Q’ is true based on the belief that it is a TT conditional should not assert the conditional on this basis, for that person is in a position to state something stronger and no less brief, namely ‘P and Q’. According to this response, the conditionals (1)-(7) seem false only because they are unassertable, and their unassertability is explained by the Assert the Stronger principle. The Material Implication Account has no problem in accommodating (P2), and so this response presents a total defense of the Material Implication Account in light of the foregoing analysis.

It is possible that this sort of response represents a kind of confusion between assertability in general (which a sentence has just in case it is acceptable—that is, fit for rational endorsement) and assertability in a context (which is constrained by pragmatic features such as conversational

relevance, fit with Grice's conversational maxims, etc.). If the Assert the Stronger principle claims that these conditionals are unassertable in general, then it follows as a matter of definition that they are unacceptable. Since the employment of the Assert the Stronger principle in this context is meant to contradict the claim that these conditionals are unacceptable (which would be entailed by their falsity), claiming that they are unassertable in general will not help defend the Material Implication Account. On the other hand, the claim may be that the conditionals are unassertable in many contexts, due to pragmatic features of what is merely suggested by the conditional. However, this kind of unassertability is irrelevant. All along, we have been judging the assertability *in general* of the example conditionals. That they are unassertable in many contexts does not undermine the claim that they are unassertable in general, because it follows from it. However, most likely this objection is not merely that the example conditionals are unassertable in some respect, but that we who endorse the claims of this chapter are confusing unassertability in a context with unassertability in general and concluding that the conditionals are unacceptable, when the fact that they are unassertable in many contexts (due to being weaker than the relevant conjunctions) does not support this claim.

Frank Jackson (1979) presents several objections to the Assert the Stronger principle (pp. 112-115), which I will not rehearse here. In addition to objecting to the principle itself, as Jackson does, one's response to this rejoinder could rest on the relative strength of the phenomena: to me, the phenomena strongly suggest not only the non-assertibility in many contexts but the falsity of (1)-(7), the failure to express a true conditional. The conditionals are not just unlikely to be assertable in very many contexts; they are unacceptable, unassertable in general, and (I think) false. For my part, I can report that I am not confusing unassertability in many contexts with unassertability in general. I will say nothing more to argue against the Material Implication Account here, because I think it fails on other grounds. (See chapters One and Two.) Someone for whom the Material Implication

Account is a live option is unlikely to be convinced otherwise by the main argument in this paper, and perhaps this is as it should be.

B. Select for multiple worlds.

One way for a Possible Worlds Account to accommodate (P1) is (i) to say that even in cases in which the antecedent is true in the actual world, the selection function selects some set of worlds larger than the actual world, and then (ii) to require that the consequent be true in all (or perhaps most of) the selected worlds in order for the conditional to be true. This feature could have the result that not all TT conditionals are true. For example, even if it is true at the actual world that grass is green and Obama is President, this account could hold that the conditional (1) ‘If grass is green, then Obama is President’ is false at the actual world through the following reasoning: the selection function from the actual world to the closest set of all grass-is-green worlds selects a set of worlds larger than, but presumably including, the actual world. Since Obama is not President at all (or most) of the worlds in the selected set, the conditional (1) is false. In virtue of what is it the case that not all the nearest grass-is-green worlds are Obama-is-President worlds? There is one modification to a Possible Worlds Account that could be used here to motivate the selection of multiple worlds in the described manner. In responding to a problem that is distinct from (but related to) the problem with accommodating (P1),¹⁴ Stalnaker (1984) suggests that perhaps some measures of similarity count not at all in determining overall closeness of worlds (as does Nolan (2003), pp. 216-217). If some measures of similarity do not count in determining closeness of worlds, then the closest set of possible worlds in which the antecedent is true will contain worlds that differ with respect to some of the features that do not count at all towards determining similarity of worlds, resulting in a tie for closest possible world even in cases in which the antecedent

¹⁴ Namely, Pavel Tichy’s (1976) hat problem, as discussed in Stalnaker (1984), pp. 127-128.

is true at the actual world. Since the presidency of Obama is seemingly unrelated to the greenness of grass, it is a decent candidate for a feature whose similarity counts not at all in determining closeness to the actual grass-is-green world. This response, if successful, allows the proponent of the Possible Worlds Account to accommodate (P1) by allowing for false TT conditionals.

There are two problems with this response. First, it seems somewhat at odds with the motivation behind the Possible Worlds Account, and therefore *ad hoc*—despite its being motivated by a challenge to the Possible Worlds Account that is distinct from the challenge of (P1). The intuitive appeal of the Possible Worlds Account stems from the notion that the proper way to evaluate a conditional is to think about the world as if the antecedent of the conditional were true, making as few other changes as necessary, and then see whether the consequent is also true in the world so conceived. The suggestion that some features do not count at all in determining similarity fits ill with this motivation. The original idea is to change as little as possible from the actual world when considering conditionals, starting the evaluation at the actual world and moving outward as needed. The new idea says that features of the actual world may be changed, as long as they are somewhat neutral or unrelated to the content of the antecedent. The presumption of changing as little as necessary is lost. It is not necessary to change who holds the office of the President in order to evaluate a world like ours in which grass is green. Thus, thinking of some features of the world as not counting toward (dis)similarity seems to get the priority backwards, by focusing on the largest set of worlds that is relevantly similar rather than on the world(s) most similar in an intuitive sense to the actual world. Of course, perhaps one could abandon this initial motivation in light of the apparent need to change the account to accommodate for false TT conditionals, and this result may be worth abandoning the initial motivation.

A second problem with this response is that so little is said about the selection function that it is not clear that such a response will accommodate both (P1) and (P2). According to (P2), the

truth values of the component parts of the conditional at the actual world do seem to determine the truth of the conditional. If whoever happens to be President at the actual world does not factor into the evaluation of (1), is it also the case that whatever number happens to be rolled at the actual world does not factor into the evaluation of (9), 'If Sophia rolls the die, the die will come up six'? Stalnaker and Nolan say so little about which features of the world count towards similarity that it is not clear that straightforward true conditional predictions will be counted as true, whereas conditionals such as (1)-(7) will be counted as false. In response to this problem, a proponent of the Possible Worlds Account could add various features to try to end up with a theory that accommodates both (P1) and (P2). Perhaps, for example, considerations of consequence at the actual world could factor heavily into determining the appropriate selection function. Such an account would indeed be an improvement insofar as it fits better with the data. In fact, the Possible Worlds Account thus supplemented might end up giving the same truth value assignments as my own account (Chapter Five), as long as it does not radically relativize conditionals to the epistemic statuses of their utterers. In such a case, there would be little to decide between the accounts besides, perhaps, that one relies on the use of an esoteric possible worlds model that the other does not need to make use of.

C. Distance the account from the actual world.

A third potential means of accommodating the competing phenomena of (P1) and (P2), and a second means of vindicating the Possible Worlds Account in particular, involves distancing the account from the actual world. Stalnaker's (1984) account includes a feature distancing it from the actual world for independent reasons. In Stalnaker's Possible Worlds Account, the selection function operates on the firm beliefs of the person considering or uttering the conditional rather than on the actual world as such. Thus, the truth that is ascribed to a true conditional may not be truth at the

actual world, since the world selected from is not the actual world. For example, a person who does not know that Obama is President will not have a selection function according to which the nearest 'grass is green' world is also an Obama-is-President world, and so (1) will come out false for such a person. Assuming that most people's belief sets are not coextensive with the set of all truths about the actual world, this account has the result that not all TT conditionals are true.

There are, again, two problems with this kind of approach. First, this approach is objectionably disconnected from the actual world. Similarly to Suppositional/Probabilistic Accounts, this account suffers from Gallimore's problem, as I will discuss in Chapter Four. According to such an account, if a person x does not know that a helium-filled balloon will rise, as a matter of the laws of physics, then it is false for x that if x lets go of a helium balloon, that balloon will rise. I direct the reader to Chapter Four for a fuller discussion of these issues, as well as citations.

The second problem, which again is shared by Suppositional/Probabilistic accounts, is that accounts such as this one escape the problem of judging all TT conditionals to be true only at the expense of judging most BB conditionals to be acceptable, which is likewise shown to be false by my examples (1)-(7). If the set of firm beliefs of the person uttering or evaluating the conditional determines which worlds are selected as closest, then the person's beliefs about the world will render some conditionals true whose component parts are irrelevant, etc., unless the account also includes some feature such as that discussed in the previous section to ensure that multiple worlds are selected for in cases of unacceptable TT conditionals. Thus the proponent of the Possible Worlds Account who relies on distancing the account from the actual world is stuck between judging most BB conditionals to be true or else adding a feature that is of questionable helpfulness.

IV. Conclusion

The major accounts of indicative conditionals all get it wrong with respect to at least one of our two phenomena: that there are false 'TT' conditionals and that some conditionals seem to be such that their 'TT' status renders them true. Such a discussion paves the way for a re-introduction of a kind of account that has gotten little attention—and none of it laudatory—since Strawson introduced it in the 1980s. Since my own account assigns truth values to indicative conditionals, I discuss the prospects for assigning truth values to indicative conditionals in the following chapter, and then finally I present and defend my own account in Chapter Five: the Consequence Account of indicative conditionals.

CHAPTER FOUR

IN FAVOR OF TRUTH VALUES

Abstract: There is strong disagreement about whether indicative conditionals have truth values. In this chapter, I present some pre-theoretical reasons for assigning truth values to conditionals. I then address four arguments that conclude that indicative conditionals lack truth values, showing them to be inadequate. Finally, I present further benefits to having a “worldly” view of conditionals, which tends to go along with assigning truth values to them.

Disagreement about the correct account of indicative conditionals involves disagreement about whether or not indicative conditionals have truth values.¹ Linguistic items that have truth values are contrasted with linguistic items such as questions, interjections, and expletives, to which truth and falsity do not apply. While there are pre-theoretic reasons for thinking that indicative conditionals have truth values, there are well-known arguments to the contrary. In this chapter, I present several initial reasons for thinking that conditionals have truth values (section I). I then respond to four arguments that conclude that indicative conditionals lack truth values. The arguments of Lewis (1976) and Edgington (1986) both fail by unfairly ruling out various truth value assignments for the whole conditional, based on the truth values of its component parts (sections II and III). Gibbard (1981) and Barnett (2006, 2012), on the other hand, both fail to recognize some good reasons for thinking that conditionals they claim to be true or acceptable are in fact false (sections IV and V). Finally, I discuss the link between an account’s ascribing truth values to indicative conditionals and its being rooted in the extra-mental world and explore some of the

¹ For simplicity, I assume throughout that true and false are the only truth values, but nothing in this chapter depends on this assumption. What I have to say applies equally well in the case of many-valued logics, because the question at issue is whether indicative conditionals are apt for the kind of properties that truth and falsity are, and the other values (if any exist) are values of the same kind.

benefits of the kind of account that has both these features (section VI). My aim in this chapter is not to establish definitively that indicative conditionals have truth values, but rather to show that the arguments to the contrary are in some way unacceptable and that there are important benefits to ascribing truth values to indicative conditionals. In doing so, I make way for my own account of indicative conditionals, which I present in Chapter Five.

I. Initial reasons for thinking conditionals have truth values

There are several reasons for thinking that indicative conditionals have truth values. The following reasons show that we seem to treat conditionals as though they have truth values, and changes would need to be made to the way we treat conditionals if it turned out that they lacked them.

A. The speech acts in which we use conditionals

One initial reason for thinking that indicative conditionals have truth values is that people use conditionals in speech acts that seem to presuppose that they have truth values. People routinely assert, deny, attack, defend, agree with, and disagree with conditionals. People who disagree about applied ethics might disagree about what to do *if* you are in a position to save people's lives by lying. People who disagree about an afterlife disagree about what will happen *if* they die tomorrow. People who disagree about the rules of tennis might disagree about who gets the ball *if* the ball hits an outer line of the court. The ability to use conditionals for these speech acts (asserting, etc.) is *prima facie* reason for thinking that conditionals have truth values, because these speech acts seem to be equivalent to expressing an attitude about the truth value of a proposition. Asserting P seems equivalent to asserting the truth of P. Denying P seems equivalent to denying the truth of P. Defending P seems equivalent to defending the truth of P. Etc. This evidence is only *prima facie*,

because there are certain locutions that likewise seem able to be asserted, denied, etc., and that some think are best analyzed as expressions of preferences. For example, I might sincerely utter, ‘This cake is delicious’, and what seems to be an assertion about a property of this cake might rather be an expression of my own response to the cake. Perhaps someone who says, ‘Pickles are delicious’ does not disagree with someone who says, ‘Pickles are disgusting’, just as someone who says, ‘Red is my favorite color’ does not disagree with someone else who says ‘Blue is my favorite color’. The two merely express different preferences. For this reason, the use of indicative conditionals in speech acts of asserting, denying, etc., counts only as *prima facie* evidence (and not proof) that conditionals have truth values—but it counts nonetheless. A second reason why this evidence is only *prima facie* evidence is that it may be that one asserts something when sincerely uttering a conditional, even if conditionals lack truth values. According to Barnett (2006), indicative conditionals lack truth values, and to assert a conditional is to assert the consequent in the context of supposing the antecedent. His account does away with truth values for the conditional itself while maintaining that something is asserted when one sincerely utters a conditional.

B. The use of conditionals in logic

Another reason for thinking that indicative conditionals have truth values is that they play a role in logic, which seems to deal with only those statements that are truth-valued. Students in introductory logic classes are taught to formalize arguments by eliminating parts of natural language arguments such as interjections and transition words that are inexpressible in terms of truth value, and conditionals are not so eliminated. At a first pass, logic seems to involve only statements that are true or false. In fact, validity is often defined in terms of preservation of truth from the premises to the conclusion, and soundness is validity plus all the premises being true.

Alternate definitions of validity and soundness are available and have been proposed by philosophers who deny that indicative conditionals have truth values. For example, Adams (1975) presents what he calls the probabilistic soundness criterion for arguments, according to which an argument is probabilistically sound just in case it is impossible for its premises to be probable while its conclusion is improbable. (Because Adams's criterion does not require that the premises themselves be highly probable [or more probable than not], it seems to be a better replacement for classical validity than for soundness.) This criterion is meant to obviate the need for premises with truth values, so probability here is not probability of truth, but some other conception—unless one construes probability of truth itself as sometimes being conditional probability of truth. Edgington (1986) seems to endorse Adams's criterion. (See pp. 29-30.) So, changes can be made to accommodate the claim that indicative conditionals lack truth values. Nonetheless, the fact that classical logic deals with conditionals and also deals only with truth-valued statements counts as evidence that indicative conditionals have truth values. The accommodating changes would be necessary only because logic is otherwise treated as dealing only with truth-valued statements, which is good initial evidence that conditionals are among these kinds of statement.

C. Embedding

A third reason for thinking that indicative conditionals have truth values is that conditionals can be unproblematically embedded in truth-functional locutions such as conjunctions. A conjunction 'A and B' is true just in case A is true and B is true. It seems perfectly germane to replace A with 'If I go to the interview, I will wear dress shoes', and B with 'If I go to the beach, I will wear sandals', and the conjunction is still true just in case the conjuncts are both true. The ease of embedding conditionals in truth-functional locutions is best accommodated by a theory according to which conditionals have truth values. If conditionals lack truth values, it will be unclear how to

treat conjunctions some of whose conjuncts are conditionals—and similarly with disjunction, negation, etc. In fact, the difficulty of accommodating non-truth-valued conditionals embedded in truth-functional locutions is the reason Lewis (1976) cites for preferring the Material Implication Account to Adams’s account (p. 305). Again, this is not to say that no accommodations can be made to the treatment of embedded conditionals, should it turn out that conditionals lack truth values. Vann McGee (1989) presents one example of an attempt to make such an accommodation. Rather, I intend to point out that the way people do treat conditionals is as though they had truth values. Thus, the embedding of conditionals in truth-functional locutions is another piece of initial evidence that conditionals have truth values.

In the following sections, I defend the claim that indicative conditionals have truth values against arguments given by Lewis (1976), Edgington (1986), Gibbard (1981), and Barnett (2006, 2012). The aim in these sections is to erode some of the argumentation that has been built up against the claim that indicative conditionals have truth values, in order to defend my presentation of an account according to which they do.

II. Argument One: Lewis’s Triviality Result and the Role of Conditional Probabilities

A. Lewis’s Triviality Result

In a context in which the Material Implication Account and Stalnaker’s Possible Worlds Account dominated the conditionals landscape, David Lewis (1976) presented an argument that would undermine both.² Lewis’s argument employs the following assumption:³

Lewis’s Assumption: The probability of a conditional ‘If A, then B’ is equal to the conditional probability of B, given A.

² Lewis actually presents two arguments, but since the first contains unnecessarily stronger assumptions than the second, and the two have the same conclusion, I will present only the second of the two arguments.

³ In Chapter Three, I referred to this claim as ‘Adams’s Hypothesis’, but I rename it here.

(Here and throughout the chapter, probability is to be construed as subjective probability.)

However, Lewis uses this assumption only in order to reject it. The result of his argument (known as Lewis's triviality result), is that it is false, given certain assumptions, that for every two propositions A and B, there is a third proposition C whose probability is equal to the conditional probability of B, given A. In other words, Lewis's Assumption can be true only if the conditional does not express a proposition; for, when it is treated as a proposition with the specified probability, it entails a falsehood. Lewis's Assumption, in addition to being initially plausible, was widely endorsed in some form or another at the time at which he published his triviality result. Lewis cites Jeffrey (1964), Ellis (1969), and Stalnaker (1970) as some early employers of the hypothesis. Ernest Adams (1965, 1975) also endorses a version of Lewis's Assumption, but his version governs the assertability of a conditional rather than the probability of its truth.

Lewis's argument proceeds in two stages. First, Lewis uses standard probability axioms to show that the probability of 'If A, then B' (construed as the conditional probability of B, given A) should be equal to the probability of B. It should be immediately clear that this subconclusion is bad news. This subconclusion is tantamount to saying that in any conditional 'If A, then B', A and B are probabilistically independent—that is, that the truth or falsity of A has no effect on the likelihood of B. But the antecedent and consequent of a conditional are not always probabilistically independent, and in fact conditionals are most often used in cases in which the occurrence of A does or would greatly affect the probability of B, usually by raising it. For example, it would be bad news if the probability of 'If you take a ride on the spaceship, you will be in outer space tomorrow' were equal to the probability that you will be in outer space tomorrow regardless of whether you take a ride on the spaceship. Contrary to the subconclusion at which Lewis arrives in stage one, the probability of B, given A, is not always equal to the probability of B.

Proving the falsity of this subconclusion is the second stage of Lewis's argument. Those who are already convinced of its falsity can skip the proof in this paragraph, which is somewhat technical. In this stage, Lewis asks us to consider a case in which there are three sentences C, D, and E, which are pairwise incompatible, but each of which has a positive probability. Suppose, for example (mine, not Lewis's), that you are drawing a single ball out of an urn filled with crimson-, daffodil-, and ecru-colored balls. Let C, D, and E be sentences corresponding to your having drawn, respectively, a ball of one of those three colors. Let A in the conditional 'If A, then C' be equal to $C \vee D$. In other words, we are considering the conditional, 'If I draw a crimson-colored ball or I draw a daffodil-colored ball, then I will draw a crimson ball'. In this case, the antecedent A and consequent C are not probabilistically independent, and the probability of C, given A, is not equal to the probability of C. (I assume throughout that one's subjective credences here are appropriately sensitive to the objective probabilities of drawing one color ball or another.) The probability of C, given A, is higher in this case than the probability of C, because the proportion of crimson balls to crimson-or-daffodil balls (the probability of C, given A) is higher than the proportion of crimson balls to crimson-or-daffodil-or-ecru balls (the probability of C on its own). Thus, the equivalence reached in the first stage of the argument does not hold for English conditionals; the present case is a counterexample. It is false in general that the probability of a conditional is equal to the probability of its consequent. This equivalence would only hold of a language that did not allow for three or more pairwise incompatible sentences all to have positive probability—a language that, unlike English, is able to express only trivial probability functions.

Since Lewis reached this unpalatable result with the assumption that the probability of 'If A, then B' is equal to the probability of B, given A, one way to avoid this result is to reject Lewis's Assumption. Another is to reject the standard probability axioms on which he relies—or else deny that they apply straightforwardly to indicative conditionals.

B. Rejecting Lewis's Assumption

For our purposes, the details of Lewis's argument are not important, because my response to them is to reject Lewis's Assumption, on independent grounds.⁴ The reason I reject Lewis's Assumption is that there are two kinds of conditionals that it fails to account for. First, I will present these two kinds of conditionals and show how Lewis's Assumption fails to accommodate them, and then I will explain why Lewis's Assumption was initially plausible in spite of these failings and what place it has in a good account of indicative conditionals.

The first kind of conditional that Lewis's Assumption fails to account for is a false (or unacceptable) 'TT' conditional. I will not here repeat my reasoning in favor of thinking that there are some unacceptable 'TT' conditionals, which I discuss at length in Chapter Three. Because Lewis's Assumption deals with subjective probability, it is prudent to speak of what we can call a BB conditional: a conditional whose antecedent and consequent are both believed, i.e., a conditional both of whose component clauses enjoy a high enough subjective probability to be believed and, we may add, asserted. (Not all conditionals where the utterer has a high degree of confidence in the antecedent and the consequent will qualify as BB conditionals, since people do not always believe propositions in which they have a high degree of confidence, as lottery cases show. Here I am dealing only with the case in which the antecedent and consequent are believed enough to be assertable for the utterer.) Lewis's Assumption entails that all high-conditional-probability BB conditionals are assertable. The conditionals I use as examples of 'TT' conditionals are also high-conditional-probability BB conditionals. Thus, if I have shown in Chapter Three that these conditionals are unacceptable, I have given a counterexample to Lewis's Assumption.

⁴ For discussions of the more formal aspects of Lewis's argument, see, e.g., Hajek and Hall (1994).

The second kind of conditional for which Lewis's Assumption fails to account is conditionals whose antecedents have zero probability.⁵ The conditional probability of B, given A, is equal to the probability of A and B, divided by the probability of A. If the antecedent has zero probability, then the denominator of the fraction is zero, in which case the value for the conditional probability is undefined. The fact that Lewis's Assumption accords an undefined conditional probability to conditionals whose antecedents have zero probability is well-noted. This fact is met with acceptance by some. For example, Edgington (1986) defines the indicative conditional as the conditional whose antecedent is an epistemic possibility, thus excluding the possibility of an indicative conditional with a zero-probability antecedent at the outset. McDermott (1996) likewise embraces the result. McDermott draws inspiration from bets, and since a conditional bet is called off if the event in the antecedent does not obtain, he is quite content to treat a conditional with a zero-probability antecedent as having an undefined truth value. Others are less enthusiastic about the result and make stipulations to avoid it. For example, Stalnaker (1970) stipulates that conditionals with zero-probability antecedents are trivially true. Some balk at discussions of zero-probability propositions, because they are inclined to treat all propositions except for very special cases as having at least some—perhaps vanishingly small—subjective probability. The idea is that perhaps I am wrong about arithmetic, seemingly simple entailment relations, and other subjects on which I am fairly certain, and so I should assign credences in these areas to reflect that I am aware that I could be mistaken about them. In reality, I think one's assignment of credences to claims in this vicinity varies somewhat depending on how careful one is being—for in other contexts people are happy to assign 0 credence to a proposition they very strongly believe to be false, even though they know that it is possible that they are mistaken. (For example, Edgington makes such an assignment in her

⁵ Brian Weatherson (2001) also provides a good reason for rejecting the claim that all conditionals with zero-probability antecedents are false, namely, that "it is a platitude that $p \rightarrow p$ is true for every p " (p. 212, with ' \rightarrow ' as the symbol for 'if...then').

objection to the Consequence Account, Chapter Five, section III.C.) Thus, it is worth discussing the merits of allowing for conditionals to have zero-probability antecedents. Those for whom the following examples do not have zero-probability antecedents will no doubt be unconvinced.

The reason conditionals with zero-probability antecedents need to be accounted for is the same reason they should not all be treated as trivially true—namely, that they can be true or false, as well as informative. Consider the following conditionals with zero-probability antecedents:

- (1) If I do not exist, then I do not have an M.A. in philosophy. [true]
- (2) If 3 is greater than 4, then I misunderstand math. [true; possibly informative to someone who does not know how to count]
- (3) If the moon is made of milk, then nothing is made of milk. [false]

Blackburn (1986) gives an example of a conditional with a zero-probability antecedent that might be useful, though his interest is in the consequent: ‘If I put my hand on this stove I will burn it’ (p. 222). In fact, there are cases in which discussions of indicative conditionals with zero-probability antecedents can be philosophically useful. For example, suppose someone is convinced by a version of the logical problem of evil, according to which the evil that occurs in our world is logically incompatible with the existence of God. Suppose also that this person is certain that evil exists (she has experienced some evil herself), and so is certain that God does not exist. This person can meaningfully consider the conditional (4), and, I think, should accept it:

- (4) If God exists, then there is some morally sufficient explanation for the evil that actually occurs.

Such a person might reasonably accept (4) because she also accepts the following: if God exists, then I am mistaken about the logical implications of evil. The person who starts out by thinking that the existence of God and the existence of evil are incompatible has one of two choices when faced with the supposition that God exists: either there is not actually evil in the world (as one had previously thought) or else one had been wrong about the logical incompatibility of God and evil. (Note the difference between the indicative (4) and the counterfactual, ‘If God were to exist, there would be

some morally sufficient explanation for the evil that actually occurs'. One can accept (4) and reject this counterfactual, because of the belief that, if God *were* to exist, then there would be no evil. This belief is compatible with the belief that, if God *does* exist, there is some morally sufficient explanation for the evil that actually occurs.)

One could respond that such a person should not place zero credence in the claim that God exists, but I think such a response is unfair. We are talking about a person who believes that the existence of God and the existence evil are *logically* incompatible, and furthermore has personally experienced evil in a way that is accessible through undeniable introspection. We can even add that the conception of evil in play here is such that even a false memory of some painful experience would itself count as an instance of evil. It may be that the logical incompatibility of the two does not withstand scrutiny, but it is nonetheless understandable that one can be (and many have been) in this precise epistemic situation. Considering the indicative conditional might be useful for the person in this epistemic situation precisely for the purpose of showing that she oughtn't place zero credence in the existence of God, or, alternatively, for encouraging her to consider creatively some options for how one might tell a consistent story about God and an evil-containing world. Additionally, considering this conditional might help this person to understand why someone who does believe that there is a God would try to tell a story according to which there is a morally sufficient reason for God's allowing evil—namely, because if God exists, then (unless there is not actually any evil) there has to be such a reason. All this is to show that there are plausible examples of conditionals to whose antecedents some people accord zero subjective probability, which nonetheless are informative, true, and even philosophically interesting. Treating all such conditionals as undefined is unacceptable, and stipulating them to be trivially true fits ill with the fact that some of them are interesting, in addition to missing the truth value on false instances such as (3) above.

C. Why Lewis's Assumption was so initially plausible and where it fits in a good account

I have presented two kinds of conditionals that are not well accounted for by Lewis's Assumption, and which therefore serve as reasons for rejecting it—and let us not forget that Lewis himself introduced the assumption in order to reject it based on the result to which it led. What, then, accounts for the initial plausibility and wide acceptance of the equivalence of a conditional's probability with the conditional probability of the consequent, given the antecedent? One reason for the initial plausibility of Lewis's Assumption is that it mirrors the decision procedure people use when deciding whether to accept a conditional. The Ramsey test that Adams's hypothesis stems from is psychologically plausible. It accurately reflects what people do when they consider a conditional—namely, hypothetically add the conditional to their stock of beliefs and consider whether or not the consequent is true in that case.

Another reason why Lewis's Assumption is so plausible (and why people may not have noticed the problems with the above two kinds of conditional) is that most of the conditionals that people have any reason to consider do not fall into either of the afore-mentioned categories of counterexamples. We can begin with false TT (or unacceptable BB) conditionals. As discussed in Chapter Three, there are two kinds of false TT conditional: (i) conditionals whose antecedents and consequents are irrelevant to each other and (ii) "Gettier-style" conditionals whose parts, though relevant to each other, are both true because of a twist or two of fate. We seldom have any reason to consider conditionals whose parts are irrelevant to each other, because they are more or less useless in guiding future action or reasoning. We use conditionals most often to discuss the connections between events or ideas, and these kinds of conditionals by definition lack such a connection. Nonetheless, they are well-formed conditionals that any account of indicative conditionals should accommodate. Moving on to the second kind of false TT conditionals: Gettier-style conditionals. We seldom consider conditionals under this particular kind of unusual circumstance because, well, it

is an unusual kind of circumstance! The cases used to show that there can be false 'TT' conditionals whose parts are relevant to each other only turn out to be 'TT' conditionals because of some unusual (though physically possible) confluence of events. It makes sense that we would not often consider this kind of conditional because of the peculiarity of the circumstances that make its parts true.

Finally, we do not often consider conditionals whose antecedents have zero subjective probability for us. If a person's credence in the antecedent of a conditional is zero, then it might seem pointless to discuss what follows from that antecedent. If you are absolutely sure that you will not attend the gala (suppose you are in a space shuttle on Mars, and the gala begins in an hour on Earth), then it is pointless to consider whether or not, if you *do* attend the gala, you will wear the red dress. (This situation is in slight contrast with the situation of considering the counterfactual 'If you were to attend the gala, you would wear the red dress'. Such a thing might wistfully, if somewhat idly, be considered.) Indicative conditionals are most often considered in cases in which they might be useful in guiding a person's actions or reasoning, and they tend to lack this usefulness in cases in which the considerer is sure that the antecedent is false. The fact that they are seldom useful helps account for the initial plausibility of Lewis's assumption; the cases that prove problematic for the assumption are not cases one would encounter regularly.

Now that we have seen why Lewis's Assumption is false and where its initial plausibility comes from, we are in a position to see what the proper role of conditional probability is within an account of indicative conditionals. A version of Lewis's Assumption can be presented as a good guide to the assertability of a conditional: a conditional 'If A, then B' is assertable only if it has a high conditional probability—that is, only if the probability of B, given A, is high—*unless* the reason it does not have a high conditional probability is that the probability of the antecedent is zero. (Here I have in mind what one might call "assertability in general," which is similar to acceptability, rather than assertability in some particular context, for which there will be a great many additional

conditions that are necessary for assertability. For more detailed discussion of assertability in general, see Chapter Three: TT Conditionals, section I.C.) The problem with Lewis's treatment of conditional probability—as well as that of Adams (1975), who treats the conditional probability as the sole determiner of assertability—is that having a high conditional probability, while being a good guide to a conditional's assertability, is neither necessary nor sufficient for it. I argue above that there are cases in which the conditional probability is high while the conditional is false. These cases show that a high conditional probability is not sufficient for an assertable conditional, because a false conditional is unassertable. The above cases in which a conditional whose antecedent has zero subjective probability is nonetheless assertable show that it is not necessary either.

III. Argument Two: Edgington's Dilemma

A second argument against truth values for indicative conditionals is presented by Dorothy Edgington (1986). Edgington's argument sometimes employs reasoning consistent with Lewis's Assumption, and so it should come as no surprise to find that it likewise faces problems with respect to unfairly ruling out truth values for the whole conditional based on the truth values of the antecedent and consequent. The argument is presented as a dilemma: if indicative conditionals have truth conditions, then either they are truth-functional or they are non-truth-functional. Edgington argues that, if they are truth-functional, then the truth conditions are those of the Material Implication Account, which are unacceptable. Edgington then presents a long argument against the other option: non-truth-functional truth conditions. She concludes that the indicative conditional lacks truth conditions, and so lacks truth values.

Since I agree with Edgington that the Material Implication Account fails, I will not rehearse her objections here. Rather, I will present and respond to Edgington's argument against non-truth-functional truth conditions. Edgington's argument against non-truth-functional truth conditions has

four parts, one for every possible combination of truth value assignments to the component parts of the conditional: TT, TF, FT, and FF. Edgington points out that if conditionals have non-truth-functional truth conditions, then there must be at least one row of the truth table for which there are multiple options, given the truth values of the antecedent and consequent, for the truth value of the whole conditional. (I will refer to such a state as the conditional being ‘open’—i.e., possibly true and possibly false.) Otherwise, the truth values of the antecedent and consequent would be sufficient for the truth value of the whole conditional, in which case the account would be truth-functional.

Edgington proceeds, row by row, to show that giving multiple options for the truth value of the whole, given the truth values of the parts, is counterintuitive. Below I present objections to three of Edgington’s four subarguments, arguing for openness of three of the rows of the truth table for the indicative conditional. Because all that is required for a non-truth-functional truth-conditional account is that conditionals at *at least one* row of the truth table are open, my rejoinder to Edgington includes much more than would be needed to establish the possibility of a good non-truth-functional truth-conditional account. Thus, even if only one of my objections below succeeds, I have succeeded in answering Edgington’s challenge. Note that, since Edgington is operating under the assumption that conditionals have truth values, in order to motivate the dilemma, I will not finesse the following to speak of acceptability/unacceptability in place of truth/falsity. My claims about the truth and falsity of particular conditionals fall under Edgington’s assumption, and so they are not instances of begging the question against the conclusion that conditionals lack truth values.

A. Edgington on TT Conditionals

Edgington begins with TT conditionals: conditionals whose antecedents and consequents are both true. Her objection to treating TT conditionals as open is that it is contrary to her criterion for an account of conditionals, which is that one’s judgment of the conditional goes by the

(subjective) conditional probability one assigns to it. Edgington also presents her criterion in other terms, which she sees as equivalent to this version, with the caveat that this version assumes that a precise numerical value can be attached to credences. Edgington seems to want to remain neutral on that assumption, but it is useful to consider this version because of its relation to Lewis's Assumption, discussed at length above. The main difference between Lewis's Assumption and Edgington's criterion is that her version has to do with whether or not a person *accepts* a conditional (finds it fit for rational endorsement) rather than an assignment of probability of truth to the conditional. Thus, Edgington's criterion is similar to Adams's (1975) treatment of conditional probabilities.

Edgington's objection to treating TT conditionals as open is that doing so violates her criterion, and it certainly does. As I argue above, the claim that some TT conditionals are false or unacceptable is incompatible with the claim that their acceptability goes by their conditional probability. My response is that, because the intuition behind Edgington's criterion can be preserved and accounted for in an account according to which the criterion itself is rejected (again, see above), the best way to deal with this clash of intuitions is to reject Edgington's criterion and accept that TT conditionals are open.

B. Edgington on TF Conditionals

For the second part of her argument, Edgington argues that there is only one possible truth value for a TF conditional (i.e., a conditional whose antecedent is true and whose consequent is false), namely, falsity. I agree with Edgington on this point, and so no further discussion is needed.

C. Edgington on FT Conditionals

The second part of Edgington's argument deals with FT conditionals—conditionals whose antecedents are false and whose consequents are true. Edgington focuses on a case in which we are certain of the consequent and uncertain of the antecedent. We are unsure whether or not a friend has mailed a letter to us, but we are sure we haven't received one. The conditional here is (5):

(5) If he mailed a letter, I didn't receive it.

Edgington's greatest complaint about this case is with how it is handled according to a Possible Worlds Account. She asks us to suppose that the friend *did* send a letter, making (5) an FT conditional. Edgington thinks it's wrong to require the truth value of 'If A, then B' here to depend on the closest possible world to a world in which the friend *did* send a letter. Of course, Edgington's criterion judges (5) to be acceptable, because the probability of the consequent, given the antecedent (using A for the antecedent here), is equal to $\Pr(A) \cdot 1 / \Pr(A)$, which is one. The fact that one is certain of the consequent is sufficient, according to Edgington, for the acceptability of the whole conditional.

I have no quarrel with Edgington's assessment of the oddness of the way the Possible Worlds Account treats conditionals such as (5). Nevertheless, Edgington has not established that all FT conditionals are true. The first problem with Edgington's argument is that cases in which one is certain of the consequent and uncertain of the antecedent are cases that are best expressed by 'even-if' sentences. See the Introduction to this dissertation, section I.B, for a discussion of these kinds of sentence and reasons for thinking that they are not equivalent to a standard indicative conditional and that a true 'even-if' sentence does not entail a true 'even'-less 'if'-sentence. The second problem with Edgington's argument is that she unfairly restricts discussion to cases in which one is uncertain of the antecedent, thus excluding consideration of FT conditionals that are clearly false. For

example, consider a case in which Brown ran in and won an election in a district in which it is not possible to win without running (no write-ins are allowed, etc.). Consider conditional (6):

(6) If Brown did not run in the election, Brown won.

In this case, (6) is an FT conditional that seems false. Given the rules of the election, it would be impossible for Brown to win without running. The fact that Brown won is not sufficient for making the truth of (6), because (6)'s antecedent is incompatible with its consequent. Even clearer cases are cases in which an antecedent contradicts a consequent. Consider (7), whose consequent is true:

(7) If grass is not green, then grass is green.

(7) is clearly false upon initial consideration. In fact, any conditional of the form 'If not-A, then A' is false, and it is a cost for any account if it treats them as true or acceptable, which it must if it treats all FT conditionals as true or acceptable. (Again, I do not beg the question against the conclusion that indicative conditionals lack truth values, because these claims fall under Edgington's assumption that conditionals have truth values.)

D. Edgington on FF Conditionals

In her argument regarding FF conditionals (conditionals both of whose component parts are false), Edgington asks us to consider a case in which we know that John and Mary will spend the evening together, but we do not know where they will spend it. She then asks us to consider a conditional such as (8):

(8) If John goes to the party, then Mary will go to the party.

Edgington builds into the case that John and Mary do not go to the party, thus making (8) an FF conditional, though of course we are considering the case in which we who evaluate (8) are meant to be unaware of this fact. Edgington holds that the only reasonable assessment of cases such as (8) is as acceptable, in which case FF conditionals are not open.

Edgington's argument here strikes me as exceedingly strange. The case in which one knows that the antecedent and consequent have the same truth value, *because* one knows that there is some established connection between the events or ideas in the conditional, seems to be a very special case. The case would read very differently if what we knew was that John did not go to the party and that Mary did not go to the party. This alone would not give us the same information about (9) that we have in Edgington's case. In Edgington's case we have an established connection between John's location and Mary's location, and I submit that it is this connection that accounts for the truth and acceptability of (9). However, such a connection is not present in all FF conditionals. Some FF conditionals are true, and some are false. Consider the following examples, with details filled in to establish their truth values:

True FF Conditionals

- | | |
|---|----------------------------|
| (9) If you eat the poison, you'll get sick. | [You do neither.] |
| (10) If I drop the glass, it will shatter. | [I don't, and it doesn't.] |
| (11) If it's 105 degrees outside, it's very warm out. | [It's neither.] |

False FF Conditionals

- | | |
|--|---|
| (12) If she ran in the election, she won. | [...of a very unpopular public figure who did neither.] |
| (13) If it's 105 degrees outside, it's below freezing. | [It's neither.] |

These examples show that the FF conditional is open. Some instances of it are true, and some are false.

As I have argued, Edgington's argument for the conclusion that there can be no two truth value assignments for indicative conditionals, given the truth values of their parts, fails on three counts: her assessment of TT conditionals, FT conditionals, and FF conditionals. Just as we found with respect to Lewis's Assumption in section II, the possibilities for various truth values for the whole undermine the argument that conditionals lack truth values.

IV. Argument Three: Gibbard's Sly Pete Case

The arguments presented by Allan Gibbard (1981) and David Barnett (2012; see section V) present a departure from the kind of argument discussed thus far. Both Gibbard and Barnett present cases that are meant to show that indicative conditionals lack truth values by showing that two people can reason validly to contradictory conditionals, using premises of which they are justifiably certain. Contradictory conditionals are conditionals one of which has the form 'If A, then B' and the other of which has the form 'If A, then not-B', where A and B are the same in each conditional. The case discussed in this section is sometimes referred to as Gibbard's Riverboat Case, but since the fact that it occurs on a riverboat is the least essential aspect of the story, I will follow those who refer to it as Gibbard's Sly Pete Case. Gibbard presents various versions of the case to make various points, but the version below is his last and most relevant version.

A. Gibbard's Sly Pete Case

The story is as follows:

Gibbard's Sly Pete Case: Sly Pete and Mr. Stone are playing a game of poker on a Mississippi riverboat. Sly Pete is a very shrewd poker player and also a bit of a cheat. Pete has two henchmen, Jack and Zack, who are also in the room with them. Jack sees both Pete's and Mr. Stone's poker hands and sees that Pete has the losing hand. Zack, on the other hand, sees only Mr. Stone's hand. Zack signals the contents of Mr. Stone's hand to Pete, unbeknownst to Jack. At this point, Mr. Stone grows suspicious of foul play and asks that the room be cleared. Jack and Zack leave the room. Zack, knowing that Pete is a shrewd poker player and that Pete knows the contents of Mr. Stone's hand, believes that if Pete called, he won. After all, Pete would not intentionally lose, and Pete knows whether or not his hand is better than Mr. Stone's. Jack, on the other hand, believes that if Pete called, he lost. After all, Pete has the losing hand, so his calling would result in a loss.

In this story, Jack and Zack believe contradictory conditionals. For ease, I will refer to these conditionals as (J) and (Z), as follows:

- (J) If Pete called, he lost.
- (Z) If Pete called, he won.

(J) and (Z) are in conflict with each other. It seems that at most one of the two can be true (*at most*, because I want to leave open the option that their truth values are undefined). If they were both true, this situation would violate the very plausible Principle of Conditional Non-contradiction: ‘If A, then B’ is incompatible with ‘If A, then not-B’.⁶ Yet it is entirely reasonable, given the information that each of them has, that Jack should accept (J) and Zack should accept (Z). This result alone would not violate the principle, because up to this point Jack and Zack have different information. Jack knows the contents of both players’ hands, and so he knows that a call from Pete would result in Pete’s losing the game. Zack knows only that Pete knows Mr. Stone’s hand, as well as his own, and that Pete is a shrewd player who would not call if his hand were the losing hand. Zack does not know the contents of Pete’s hand. It is Gibbard’s assessment of what would happen if Jack and Zack shared their information that has the potential to create problems for an account according to which indicative conditionals have truth values. According to Gibbard, if Jack and Zack shared their information with each other, so that Jack came to know that Pete knew the contents of Mr. Stone’s hand and Zack came to know that Pete had the losing hand, neither Jack nor Zack would come to believe that they had been wrong in endorsing the conditional each previously endorsed. The reason is that each would recognize that their acceptance of (J) or (Z) was well-founded on their evidence at the time. In a sense, both Jack and Zack were right when they believed the conflicting (J) and (Z), and this assessment is made, according to Gibbard, both by us who know the all the information as well as by Jack and Zack themselves after they learn each other’s information. It is important to note that Gibbard is not equivocating about the sense in which Jack and Zack are right. One might worry that Gibbard is equivocating here between being right about x in the sense of (a) not being irrational with respect to x, given one’s information, and (b) being right

⁶ One might want to restrict the principle to cases in which the antecedent is not a contradiction. Thanks to Graeme Forbes for pointing out this possibility.

by judging x , which is true, to be true. Far from equivocating, Gibbard presents an argument from being right in sense (a), together with other premises, to being right in sense (b). Gibbard says, of our (J) and (Z),

If both these utterances express propositions, then I think we can see that both express true propositions. In the first place, both are assertable, given what their respective utterers know. [...] From this, we can see that neither is asserting anything false. For one sincerely asserts something false only when one is mistaken about something germane. In this case, neither Zack nor Jack has any relevant false beliefs. [...] Neither, then, could be sincerely uttering anything false. Each is sincere, and so each, if he is asserting a proposition at all, is asserting a true proposition. (231)

By Gibbard's assessment, the Sly Pete case involves two contradictory conditionals which, if they express propositions at all, express true propositions.

This would, indeed, be a problem for a truth-valued account, because if conditionals have objective truth values, then two conflicting conditionals such as (J) and (Z) cannot both be true in the same context without violating the very plausible Principle of Conditional Non-contradiction. And yet, says Gibbard, they are. Gibbard presents his result as a dilemma: in light of the Sly Pete Case, one must either (a) treat the propositions expressed by conditionals as radically sensitive to the epistemic states of the speakers or (b) deny that conditionals express propositions at all, which entails that they lack truth values. Gibbard's own route is (b).

B. Zack was wrong: responding to Gibbard

Gibbard's Sly Pete Case has generated a great many responses in print. One of the features of the overwhelming number of these responses that has always baffled me is the acceptance of Gibbard's assessment that Jack and Zack are both right about what happened if Pete called. Stalnaker (2005) adjusts his theory in light of Gibbard's case by taking route (a) above, and Krzyżanowska et. al. (2014) likewise treat both (J) and (Z) as true with respect to a body of evidence. Lycan (2001), though first expressing sympathies with those who reject Gibbard's conclusion, goes

on to argue that Zack was right. Jackson (1990) uses Gibbard's case to argue for a particular distinction between indicative and subjunctive conditionals.

So, let me be the almost-first to say it: Zack was wrong.⁷ It is not the case that, if Pete called, he won. We may suppose, as seems to be built into the specifics of the case, that there is a true counterfactual in the vicinity:

(14) Pete would have called (while knowing all the relevant information about everyone's hands) only if calling would lead to a win.

Or, more archaically but in a more standard formulation, if Pete were not to be in a position to win (and to know himself to be in such a position), he would not call.⁸ It does not follow from this counterfactual that if Pete called, he won. The indicative conditional (*Z*) is true only if Pete had the winning hand. The fact that Pete had a losing hand is the reason that, as some who have written on the case agree (see, e.g., Lycan (2001)), pace Gibbard, if Zack and Jack were to share their information after clearing the poker room, Zack would no longer believe that if Pete called, he won. (For a divergent view, see Stalnaker (1984), p. 113.) Yet Zack and Jack may both heartily agree to the counterfactual. Pete is reliably a shrewd player, with the desire to win. Even if Zack is wrong, there is a sense among those responding to the case that we must explain why his conditional was so acceptable to himself, why Zack was not at fault in expressing it, given his information. After all, we can presume that Zack did not have any false relevant beliefs, nor did he seem to make an error in reasoning. Whence, then, his mistake? My assessment is that Zack did make an error in reasoning: he took the indicative conditional (*Z*) to be entailed by the counterfactual conditional (14), when it is not. Perhaps Pete would not have called in the situation in which he was in, but if he did call, he lost.

⁷ Pendlebury (1989) agrees that Zack was wrong, and his assessment was a relief to encounter.

⁸ Thanks to Lycan (2001), who points out that (*Z*) has a similar flavor to what Lycan calls backtracking conditionals. See pp. 178ff. Lycan's assessment differs from mine, because he thinks that (*Z*) itself should be considered to be a backtracker, whereas I contend that Zack mistakenly accepts (*Z*) as a consequence of (14).

As I have described it, Gibbard's Sly Pete case is not a puzzle. The puzzle arises only if the rather-independently-plausible Principle of Conditional Non-contradiction seems genuinely violated. But it does not. Zack was wrong. The Sly Pete case is not a case of two true or mutually acceptable conditional contradictions. Without this feature, Gibbard's dilemma between (a) treating conditionals as expressing propositions that are highly context-dependent or else (b) dispensing with conditional propositions at all simply does not arise. In fact, I am tempted to take the Sly Pete case as being evidence in *favor* of accounts that assign objective, not-highly-context-dependent truth values to indicative conditionals. An account that denies that indicative conditionals have truth values would have to respond to the Sly Pete Case by saying that both conditionals are true or acceptable, but they are not. Given the facts about the poker hands in play, it is simply false that, if Pete called, he won, and accounts which predict this result fare better in that respect than accounts that do not.

V. Argument Four: Barnett's Conscious Being Case

David Barnett (2012) presents a case that is meant to be similar in style to Gibbard's Sly Pete Case, but more persuasive. Again, it is a situation in which we are meant to judge that two people who hold contradictory conditionals in the same scenario are both justifiably certain of those conditionals. Barnett uses world leaders in his (2012) presentation of the case, but I am going to discuss a more elegant version that Barnett has brought up in personal conversation before returning to his (2012) version. Barnett says it is reasonable for him to be fully certain of (15):

(15) If there is only one conscious being, David Barnett is it.

The reason it is reasonable for David to accept this conditional with full certainty is that he knows, perhaps indubitably, that he himself is conscious, and so supposing that only one being is conscious

leads him to accept with full certainty that he himself is that being. However, you and I (who are not David Barnett) must be certain of the contrary of (15):

(16) If there is only one conscious being, David Barnett is not it.

We must add the background assumption that I who am evaluating (16) am not identical to David Barnett. (Barnett, when reading this, must substitute a different name.) My certainty of (16) is based on my belief in my distinctness from Barnett as well as my belief in the following:

(17) If there is only one conscious being, I am it.

Since I am certain that I am a conscious being, (16) is eminently reasonable to me, just as (15) is to David Barnett, and one would not fault either of us for accepting our respective conditionals.

Moreover, on a certain conception of justified belief, if one is justifiably certain of P, then P is true.

Yet (15) and (16) are contradictory. If they are both true at the same time, then the Principle of Conditional Non-contradiction is violated, because (15) and (16) are of the form 'If A, then B' and 'If A, then not-B'. Barnett concludes that indicative conditionals lack truth values, though treating conditionals as radically-indexical truth-valued sentences has not been ruled out by this case.

Barnett's Conscious Being Case is more convincing than Gibbard's Sly Pete Case, because the two parties in the case have symmetrical evidence: neither one seems better suited to make a judgment about what would be true if the antecedent of (15) or (16) obtained. It also has the merit of not requiring any imagined scenario to set the scene. Each of us seems to be in a position in which we must choose between (15) and (16), and it is clear which one each should choose. In light of the persuasiveness of Barnett's case, it is tempting to treat each of (15) and (16) as true-with-respect-to-a-body-of-evidence: (15) is true with respect to Barnett's body of evidence—namely, evidence of his own consciousness—and similarly for (16) and everyone else's body of evidence. Perhaps the conditionals are both true, but the propositions they represent are radically indexed to the epistemic states of their respective utterers. This treatment loosens the ties between a

conditional's truth conditions and the extra-mental world. I leave discussion of the dangers of loosening these ties to section VI, in which I outline some of the benefits of a "worldly" view, and limit discussion here to the alternate conclusion that indicative conditionals lack truth values.

One problem with Barnett's Conscious Being Case—though it is not the main problem—is that it is not possible for one to rationally accept all the premises at the same time. It is unreasonable for me to believe *both* that Barnett should believe (15) *and* that, if there is only one conscious being, Barnett is not it. When I believe that Barnett should rationally accept something, I presuppose that Barnett is a(n existent) conscious being, and I know that I myself am one. I cannot both presuppose that he exists and also suppose that only one of us exists.

Nonetheless, there is a further problem with Barnett's Conscious Being Case, and that is that the notion that one should be certain of (17), on which one's certainty in (15) or (16) is founded, is based on a false dilemma. One might think that 'I am a conscious being' entails that if only one conscious being exists, I am it; but I submit that it does not. The sense that this entailment holds plausibly stems from the following sort of reasoning: We are supposing there's only one conscious being, so either I'm it or I'm not. This means that, if only one conscious being exists, it cannot follow that I'm not it. Thus, it must follow that I *am* it—and from thence we get the confidence that, if there is only one conscious being, I am it. However, this reasoning is flawed, specifically in the step from 'it cannot follow that I'm not it' to 'it must follow that I *am* it'. The flaw is based on a tacit assumption of Conditional Excluded Middle, the principle according to which either 'If A, then B' or 'If A, then not-B' must be true. This kind of reasoning forces us to choose between 'If there is only one conscious being, I'm it' and 'If there's only one conscious being, I'm not it'. But if we reject Conditional Excluded Middle, we can choose to reject both conditionals. All we need is a plausible reason for thinking that neither conditional is true, and that reason is that, supposing there is only one conscious being, plausibly there is nothing in the world favoring its being me over its being

David Barnett. As Barnett himself says, “there is no objective fact of the matter whether this [consequent] is correct relative to the supposition” (p. 424). If the fact that the lone conscious being is me does not follow from the antecedent, and neither does the fact that it is Barnett, then perhaps both (15) and (16) are false.

Barnett’s (2012) example involved other people, Barack Obama and Hu Jintao, disagreeing about relevantly similar conditionals; perhaps such a version fares better. Yet we can see that it does not. I cannot both accept that Obama and Hu Jintao exist and accept that they should believe the contradictory conditionals (18) and (19), respectively:

(18) If only one conscious being exists, it is Barack Obama.

(19) If only one conscious being exists, it is not Barack Obama.

The reason, again, is that nothing about the world, as far as we know, settles the facts about who exists, if only one of these leaders exists. Both sentences are false. As of yet, I have not provided a reason for thinking that that the lack of the world “settling the facts” should render the conditionals false, and that is because elucidating this reason requires presenting an account linking the way the world is to the truth values of conditionals. I present such an account in Chapter Five, discussing the resources with which my own worldly account responds to Barnett’s Conscious Being Case in section II.G. For now, it is sufficient that we have seen one way in which a truth-valued account of indicative conditionals might respond to Barnett’s Conscious Being Case without radically indexing conditionals to their utterers or eliminating truth values. If such a response is possible, then there is still room for a truth-valued account of indicative conditionals.

VI. The benefits of a “worldly” view

In the first section of this chapter, I presented three initial reasons for thinking that indicative conditionals have truth values, and now I have defended that claim against four arguments that conclude the contrary. In this section, I briefly discuss some of the benefits of a “worldly” view

of conditionals: a view according to which the conditional has truth conditions that tie it to the world outside one's own beliefs. The Material Implication Account is a worldly account, as are versions of the Possible Worlds Account that do not treat the conditional as radically indexical (which rules out Stalnaker (2005) and Nolan (2003)).⁹ The Consequence Account is perhaps worldliest of all, because its truth conditions involve relations in the actual world.

Suppositional/Probabilistic Accounts such as those of Adams, Edgington, Gibbard, and Barnett, which eliminate truth values, do not closely tie conditionals to the extra-mental world.¹⁰

It seems appropriate that the truth conditions of indicative conditionals should be tied to the world outside one's own head. Consider (20).

(20) If I let go of a helium-filled balloon, it will rise.

(20) is true or acceptable not because of my own beliefs about helium-filled balloons but because of *facts* about helium-filled balloons—facts that were true before anyone had subjective credences about them. Non-worldly accounts miss this fact. Proponents of non-worldly accounts can say what our subjective credences about natural phenomena ought to be, in order to reflect the laws of nature, and under what conditions conditionals are acceptable, but they cannot bridge the gap between the laws of nature and the acceptability conditions of our own conditionals. A person whose subjective credences do not render 'the balloon will rise' probable, given that I let the balloon go, is someone for whom (20) is justifiably unacceptable, and there is no fact of the matter according to these accounts, because (20) does not express a matter of fact. One contributor to this problem is that non-worldly accounts do not have the right truth-makers (and not just because they most often do

⁹ Nolan (2003) holds a version of Possible Worlds Account that treats conditionals as radically indexical, but which maintains some worldliness by holding that it is our knowledge (rather than mere belief or firm belief) that determines which possible worlds are closest to ours. Krzyżanowska et. al. (2014) do something similar within their Suppositional/Probabilistic Account by requiring that the relevant epistemic states be knowledge states.

¹⁰ Some conditionals are explicitly about people's beliefs, such as the conditional 'If x believes A, then x probably believes B', but the relevant difference here concerns whether or not an account treats the acceptability or truth value of all conditionals as dependent just on a person's beliefs.

not assign truth values to conditionals). A truth-maker is the state of affairs, event, individual etc., which makes a true sentence or true proposition true. For example, my laptop, or the fact that my laptop exists, is one of the truth-makers for the sentence ‘At least one laptop exists’. If a non-worldly account is true, then what makes (20) true (for Stalnaker) or acceptable (for other kinds of account) is facts about the speaker’s beliefs, rather than facts about helium-filled balloons. If conditionals do not have objective truth values that depend on the world outside one’s own beliefs, then whether or not (20) is acceptable does not depend on physics but on oneself. Treating conditionals as radically indexical comes at a great cost. This feature of non-worldly accounts is unacceptable.

This complaint is related to what William Lycan (2001) refers to as Gallimore’s Problem, named for the person who raised it in conversation with Lycan. (See pp. 69-72.) Consider a conditional such as (21), which we have no earthly reason to accept:

(21) If I finish this chapter today, Norway will have an unusually early autumn in [2015].¹¹

It seems vanishingly unlikely that my finishing this chapter today will have any effect on the timing of autumn in Norway, in 2015 or any year, and I find (21) unacceptable. Gallimore’s Problem asks us to consider a case in which, unbeknownst to us, there is a lawful connection between the two events, such that my finishing this chapter today really does or would, as a matter of natural law, lead Norway to have an early autumn in 2015. If such a bizarre law exists (perhaps God is feeling mischievous), then of course (21) is true. Yet because we have no reason at all to believe (21), non-worldly accounts must treat it as unacceptable. Of course it is the case that, if we have no reason to accept (21), it should be unacceptable to us, even if true. The problem is that non-worldly accounts lack the apparatus to say that (21) is true-but-unacceptable, given our epistemic states. According to

¹¹ I changed the year such that it is in the future relative to the time at which this chapter is written.

non-worldly accounts, our credences in (21) settle the matter about it, and not even God's putting a law in place could make (21) true. As Lycan puts it, this case embarrasses any non-worldly account. Gallimore's Problem, of course, is only half the problem, and we needn't consider only bizarre cases in order to make the same point. Any conditionals about the laws of nature that we do not have the evidence or expertise to find acceptable *are* unacceptable according to non-worldly views. Conversely, any of our well-established but actually false theories are acceptable according to non-worldly views.¹² The world does not get any say in the matter, so to speak, except insofar as our experiences in the world inform our subjective credences. I will say it again: the flexibility that allows non-worldly accounts to accept both the contradictory conditionals in Gibbard's or Barnett's case comes at too high a cost. Non-worldly accounts place the acceptability of conditionals in our own heads, instead of in the world, as it should be.

VII. Conclusion

The claim that indicative conditionals have truth values has much to recommend it initially, and I have argued that the prominent arguments against truth-valued accounts are each problematic. Lewis and Edgington's arguments are problematic because they unfairly rule out the possibility of certain assignments of truth values based on the combination of truth values of the antecedent and consequent. Gibbard and Barnett's arguments are problematic for a different reason—namely, that they give the wrong truth value assignments to the specific conditionals used in their arguments. Furthermore, accounts that more closely tie the evaluation of conditionals to the extra-mental world (as non-truth-valued accounts fail to do) have certain advantages over accounts that lack such a tie. Given the ubiquity of reasons in favor of a truth-valued account and the poverty of reasons against

¹² This complaint is related to Edgington's (1995) rain dance case. Edgington points out that conditionals such as 'If we perform this rain dance, then it will rain' are justifiably acceptable for people who believe that rain dances bring rain, according to non-worldly accounts, whereas in reality they are false/unacceptable.

one, I conclude that the possibility of assigning truth values to indicative conditionals is indeed worth exploring further.

CHAPTER FIVE

THE CONSEQUENCE ACCOUNT

Abstract: Because many have mistakenly denied that indicative conditionals have truth values, most have ignored connection-type accounts of indicative conditionals. In this chapter, I present a new connection-type account, the Consequence Account, according to which a conditional ‘If P, Q’ is true just in case there is some relation by virtue of which Q is a consequence of P. I discuss the details of this account and the logical properties it assigns to conditionals, review what this account gets right, and defend it against objections.

Up to this point, we have seen reason upon reason for rejecting the major accounts of indicative conditionals. These accounts assign truth values for the wrong reasons or mistakenly do away with truth values altogether. Thus far I have been telling a story of inadequacy—of many problems with no hero to resolve them. I will name our hero now: the Consequence Account. In this chapter, I present this account and show how it gets right much of what the major accounts get wrong.

The Consequence Account is a “connection-type” account, which requires that there be some connection between the content of the antecedent and the content of the consequent in order for the conditional to be true. My version of connection-type account draws inspiration from the considerations advanced by Strawson (1981). Strawson likens linguistic constructions containing ‘if’ to linguistic constructions containing ‘so’, noting that conditionals seem to suggest a kind of ground-consequence relation between the antecedent and the consequent; the antecedent is presented as grounds for the consequent. To minimize confusion with talk of metaphysical ground, which is only a special case of the kind of relation Strawson discusses, I will speak just in terms of a ‘consequence relation’.

Connection-type accounts have been more or less ignored since Strawson (1981) proposed his version.¹ My assessment is that there are three reasons that partially account for the dearth of connection-type accounts in the last thirty years. First, Lewis's (1976) triviality result, which some took as proving that indicative conditionals lack truth values,² rendered people less likely to introduce new truth-valued accounts. (See Chapter Four: In Favor of Truth Values, section II, for a fuller discussion of Lewis and my reasons for rejecting his assumption.) Since connection-type accounts assign truth values to conditionals, they waned in popularity, and truth-valueless Suppositional/Probabilistic Accounts gained momentum. A second reason for the lack of connection-type accounts is Edgington's (1986) argument that all truth-valued accounts must be truth-functional accounts. I discuss this argument as Edgington's dilemma in Chapter Four, section III, where I likewise give my reasons for rejecting Edgington's conclusion about the problematic nature of non-truth-functional truth-value accounts. If there is no reasonable space for assigning truth values to conditionals without assigning them as a function of the truth values of the component parts, then there is no room for a connection-type account. Thus, anyone convinced by Edgington would be unlikely to pursue a connection-type account. Thirdly, there is the more general problem of how to make a connection-type account rigorous in such a way that it seems like a peer of the major accounts with their various technical apparatus. Pendlebury (1989) gives a rather compelling speech along these lines. He says,

And if one wants to be a realist about indicative conditionals, as I do, then there must be something *about the world* which legitimates our perfect thinker's dispositions to change what he accepts in response to new information. It is ultimately these worldly factors which make conditionals true or false. But what are they? In answer to this question I can merely gesture roughly in the direction of relations between properties, causal facts, laws of nature, and general patterns and regularities in the world as a whole or in relevant parts of it (and here I include the rules and

¹ One notable exception is Krzyżanowska et. al. (2014), who use the motivation for a consequence account to present what is essentially a kind of Suppositional/Probabilistic Account, in which the relevant relata in consequence relations are the antecedent together with the utterer's background knowledge. Their work builds on Douven (2008).

² Adams (1975, section 1.2), for example, embraced this conclusion.

regularities of human games and institutions). Pressed for further details, I shamefully admit that I cannot do much better than supply a few well-worn examples; So instead of giving you a hard-headed account of what it is in the real world which makes conditionals true or false, I follow the leaders in the field and try to translate my talk about perfect thinkers and their rational dispositions into a purely abstract characterization of the truth conditions of conditionals in terms of possible worlds. (186, italics in original)

Pendlebury all but laments his decision not to pursue a connection-type account, an account according to which relations between things in the world make true conditionals true. He does not see how one can make such an account rigorous, and I sympathize with his worry. In response to this worry, I believe the recent proliferation of work in the area of metaphysical dependence has opened up room for connection-type accounts—though not intentionally. There has been a kind of fatigue and disillusionment in metaphysics with formal features such as modality and a decreased confidence in their ability to account for metaphysical relations of dependence. Metaphysicians have thus moved towards accounts of dependence that treat it as primitive and unanalyzable. (See, for example, Fine (2001), Schaffer (2009), Rosen (2010), and Sider (2014).) I see such a move as opening up dialectic space for an account of indicative conditionals that makes use of dependence-type connections by legitimating discussion of such connections, in a way that has not been available until recently.

In the following sections, I present the Consequence Account (section I) and then discuss what it does right, in contrast with the major accounts (section II). Despite the absence of a prominent proponent of connection-type accounts, such accounts have remained in the background as suggestions to be dispensed with. Thus, there are objections to connection-type accounts even in the absence of a prominent proponent of them. I deal with such objections in section III.

I. The Consequence Account and the logical properties of ‘if’

According to the Consequence Account, a novel version of the connection-type view that I will develop and defend here, the indicative conditional ‘If P, Q’ means that Q is (was, will be) a consequence of P. Here are its truth conditions:

The Consequence Account: The indicative conditional ‘If P, Q’ is true just in case there is some relation by virtue of which Q is (was, will be) a consequence of P.

In general, the relation mentioned in the truth conditions (what I will call the consequence relation) will be not between linguistic items, the antecedent and consequent themselves, but between the events, facts, etc., that are described by the antecedent and consequent, and this relation will make it the case that the consequent is a consequence of the antecedent.

One might be tempted to formulate a connection-type account in terms of the truthmaker of the antecedent and the consequent. I think this formulation would be problematic in some respects. First, one would need to talk about *possible* truthmakers, since the antecedent and the consequent may very well be false while the conditional is true; and doing so may prove to be tricky. Second, it may be the case that some kinds of sentences lack truthmakers even when those sentences are true. (See, for example, Saenz (2014) for such a claim.) In general, it is best if an account relies as little as possible on controversial claims, and I do not want my account of indicative conditionals to depend on the controversial claim that all truths have truthmakers or (what may be equivalent) that all sentences have possible truthmakers. Thus, I formulate the Consequence Account in terms of *what is described by* the antecedent and the consequent—which are, probably, events or states of affairs. I refer to what is described by the component parts as their “content,” but I do not use this term in the technical way in which philosophers use it when speaking about propositions. Again, the reason is that I do not want my account to rely on settling controversies in other areas of philosophy one way or another. Another part of the reason is that, even if it turn out that propositions have abstract entities as constituents, the relations I am discussing need not have abstract entities as their rela-

For example, ‘If I am in Boulder, I am in Colorado’ is true because a relation between places (Boulder’s being in Colorado) makes the consequent a consequence of the antecedent, and the Consequence Account makes this judgment regardless of what individuates or constitutes a proposition.

The notion of a consequence relation should be somewhat intuitive, and I think it is best described by means of examples. The consequence relation that makes a conditional true can be of many forms, such as logical consequence (1), metaphysical grounding (2), rule-generated consequence (3), or causal consequence (4).

- (1) If x is greater than 3, then x is greater than any number less than three.
- (2) If the donut is shaped thus-wise, then the donut hole is shaped thus-wise.
- (3) If Leon moves his queen, he will check-mate.
- (4) If the vase is knocked off the counter, it will break.

The types of consequence exemplified in (1)-(4) are not exhaustive, and I will not attempt to provide an exhaustive list of the kinds of consequence relations that exist. Consequence can be thought of as a kind of relation that is often, but not always, asymmetric. The limiting case is conditionals of the form ‘If A , then A ’. Such conditionals are plausibly always true, since every proposition is a logical consequence of itself, and so logical consequence is not asymmetric. It is logical consequence as well as consequence generated by definition or analysis that accounts for the existence of true bi-conditionals. Because most kinds of dependence are asymmetric, most cases in which P has Q as a consequence are not also cases in which Q has P as a consequence, the exceptions being cases where P and Q are identical, P and Q are a definiens and its definiendum, P is analyzed as Q , etc. Most other types of consequence relation, however, are asymmetric, as in the case of each of (1)-(4). Thus, consequence can be thought of as the most general kind of dependence. Aside from the extreme cases mentioned above, consequence is one thing leading to another, one fact underlying another, in any of various forms that might take. Such a notion might seem simple, and it is. Consequence is the familiar notion underlying all our standard uses of ‘if’. One might wonder how such a simple notion

can be informative, but we shall see in the coming sections that it is, because this account renders different judgments from those rendered by the major accounts.

A. What's in and what's out

There are various candidates for relations that might count as consequence relations, and so some work must be done to determine which relations are consequence relations and which are not. This work will shed light on which conditionals are true and false according to the Consequence Account.

1. Unrelated events

One kind of case that can be ruled out immediately is a case in which Q follows P, but is not a part of the causal chain of which P is an earlier member. For example,

(5) If I roll a six, your cat will die tonight.

In this case, even if both the antecedent and the consequent are true, the cat's dying is not a causal consequence of my rolling a six. (If it *is* a causal consequence of my rolling a six—say, because someone has agreed to kill the cat upon my rolling a six, or through magic or something—then the Consequence Account gives quite a different judgment.) Whether or not (5) is true depends on whether or not the rolling of a six does or will have the cat's death tonight as a consequence, and this result seems appropriate. Thus, we can refine our judgments about true conditional predictions discussed in Chapter Three. Not all true conditional predictions come out as true, according to the Consequence Account. Only those whose antecedents and consequents describe two parts of a causal chain will count (or those where some other consequence relation applies). This result seems like a good one, because (5) seems false in a case in which the rolling of the die has nothing to do

with the cat's death. I return to the issue of 'TT' conditionals in light of the Consequence Account in section II.E.

At this point one might object that the judgments of the Consequence Account in this case diverge from intuitions about betting. If one were to bet that, if I roll a six, your cat will die tonight, and I *do* roll a six and your cat *does* die tonight, then it seems that the person has won the bet. But what is a bet except a contest concerning the truth of the statement bet upon? According to the Consequence Account, (5) is false, but a person who bet by means of (5) would win the bet if (5) turned out to be a 'TT' conditional. This situation shows that judgments about betting and the Consequence Account's judgments about the truth values of indicative conditionals diverge. This result should not be alarming, because betting is a system in which there are three possible outcomes: a bet is either won, lost, or called off. However, if we are using a two-valued logic, a well-formed declarative sentence can only ever be true or false. Thus, the practice of betting will not map neatly onto a two-valued system of any kind. If that is the case, what goes on when a person bets in the form of a conditional prediction? My judgment is that, when a person bets in the form of a conditional prediction, the person is not betting that a conditional is true but rather is conditionally placing a bet. The better establishes a convention according to which, if the event described in the antecedent obtains, then he or she will pay or receive some sum dependent on whether or not the event described in the consequent obtains. Just as in cases of consequence generated by the rules of a game, placing a conditional bet involves a person causing there to be a consequence relation between the content of the antecedent and the existence of the bet. However, since this consequence relation is not between the antecedent and consequent themselves, it does not have any bearing on the truth value of the bet-upon conditional.

Suppositional/Probabilistic Accounts such as that of Barnett (2006) treat bets in this manner while also giving the same treatment to all other relations to conditionals, such as asserting that if P,

then Q; fearing that if P, then Q; etc. The unity of such an approach is admirable, but the special features of a bet seem to warrant treating a bet as a special case—namely, the fact that a bet has three values of being won, lost, or called off. Pace Barnett, there seems to be no equivalent to being called off for the other attitudes and speech acts, and thus it is neither *ad hoc* nor unexpected to treat betting as a special case.

2. *P's raising the probability of Q*

Another question is whether or not relations in which P (merely) raises the probability of Q count as consequence relations. (Probability here can be subjective or objective; the result is the same.) Is P's raising the probability of Q enough to make Q a consequence of P? I think we will see that it is not. Many instances in which P raises the probability of Q are also instances of some other consequence relation—for example, cases in which P causes Q or P logically entails (contingent) Q. So, in order to test whether or not P's raising the probability of Q is itself a particular kind of consequence relation, we need to evaluate a case in which P raises the probability of Q, where P and Q do not stand in some other consequence relation to each other. If the antecedent raises the probability of the consequent, and yet the antecedent does not have the consequent as a consequence, then this kind of relation is not a kind of consequence relation. Consider a case in which you have an urn filled with marbles, some black and the rest white. For our purposes, the ratio of black to white marbles does not matter, except that there are multiple of each shade (such that drawing a black marble first will not guarantee that a white marble will be drawn second). When you draw a marble, you do not then replace it in the urn. Consider (6):

(6) If you remove a black marble first, the second marble you draw will be white.

Removing a black marble first does raise the probability that the second marble drawn will be white.

Removing the black marble shifts the ratio of black to white marbles, making it more likely than it

was previously that you will draw a white marble on your second draw. However, it is not the case that your drawing and removing a black marble first has your drawing a white marble second as a consequence, even if you do draw a white marble on your second draw. This case is an instance of the antecedent's raising the probability of the consequent, but it is not a case of the antecedent's having the consequent as a consequence, and so the antecedent's raising the probability of the consequent is not a kind of consequence relation. A conditional whose antecedent raises the probability of its consequent is true only if there is some other relation that serves as the consequence relation in that case.

3. P's raising the probability of Q to some high degree

Next we may be tempted to ask about a case in which P raises the probability of Q to some high degree. Perhaps it is not enough for the antecedent to raise the probability of the consequent; it must also make the consequent more likely than not, or overwhelmingly likely. Again, in order to test whether or not the relation of P's making Q sufficiently highly likely is a kind of consequence relation, we need to consider a case in which this relation holds and no other consequence relation holds between P and Q. Thus, we cannot consider a case in which P causes Q, or P entails Q, or P and Q enter into any other type of consequence relation. Let us consider another case in which you have an urn filled with black and white marbles. Suppose that there are five black marbles and ninety-five white marbles in the urn before your first draw. Consider (6) above once again, in light of the current scenario. Your removing a black marble does not seem to have your drawing a white marble as a consequence. There seems to be no consequence relation between the two actions. The two actions, drawing one marble and then the other, are not causally linked. In this case, even if (6) is a TT conditional, your drawing a black marble first does not have as a consequence that you draw a white marble second. It makes that second draw highly likely, but neither is a consequence of the

other. Thus, P's raising the probability of Q to a very high degree does not seem to be an instance of a consequence relation.

The present case may raise concerns about the relation between the judgments of the Consequence Account and the probabilities of conditionals, and I want to discuss such a worry. If the Consequence Account is accurate, then our judgments about the probability of the consequent, given the antecedent, will not be equal to the probability of the conditional as a whole. This topic was discussed in Chapter Four: In Favor of Truth Values, and so I will not rehearse at length the reasons for thinking that the probability of the conditional is not equal to its conditional probability. However, one might make initial judgments about the probability of a conditional that do not match up with the judgment of it as true or false according to the Consequence Account. One might think it is highly probable that (6) is true, because it seems highly probable that, if you first draw a black marble out of the urn without replacing it, you will draw a white marble second; and, indeed, the probability of the consequent, given the antecedent, is very high in this case. But the lesson we learned in Chapter Four is that treating the probability of the conditional in this way leads to problems (namely, Lewis's triviality result), and that there are independent reasons for rejecting this equivalence (namely, that it mistakenly judges all BB conditionals to be acceptable and that it cannot handle conditionals with zero-probability antecedents). While it may seem odd or unnatural to separate one's judgment about the truth value of conditionals from one's judgment about their conditional probabilities, we have well-motivated reasons for doing so, and remembering these reasons should help alleviate some of the strangeness of judging (6) to be false in this case.

4. Purely epistemic associations

P's raising the probability of Q to some high degree may be insufficient for a true conditional, but there are certain conditionals that are similarly problematic and yet which have

some instances that seem to be true—or at any rate, useful and common enough. These cases are instances of what I will call purely epistemic associations: cases in which the relation between P and Q seems to be a relation just between one’s own credences in P and Q.³ These cases do not deeply differ from cases in which P raises the subjective probability of Q, but they do come across as special cases, and so I give them their own treatment. Consider a case in which my interlocutor is trying to think of the TV show I was discussing last week.⁴ Since I discussed different TV shows on different days last week, I try to help my interlocutor narrow down which show I was discussing. My interlocutor asks, “What TV show were you talking about last week?” and I respond with (7):

(7) If you’re thinking of Wednesday, I was talking about the TV show “Parks and Recreation.”

The conditional (7) expresses a purely epistemic association between last Wednesday and my talking about a certain TV show. We can even add that there is no more general association between Wednesdays and discussions of this particular show. I am merely remembering what I talked about on Wednesday and associating that day with a discussion of a certain show.

The conditional (7) seems to be true or useful in this situation, but as we have seen in Chapter Four: In Favor of Truth Values (section VI), treating all purely epistemic associations as acceptable causes problems for an account of indicative conditionals. There are conditionals that are intuitively false but which are judged to be true by such a treatment. All that is required for an example is that the person have a set of credences which do not match up with reality in the right way—for example, someone who believes that the sun orbits the earth or that lonely Salemites have dangerous magical powers. Such a person cannot make it true that if a day has passed, then the sun

³ What I call purely epistemic associations are a subcategory of Eve Sweetser’s (1991) epistemic conditionals. Sweetser, for example, treats conditionals where the consequent is analytically entailed by the antecedent as epistemic conditionals, such as ‘If she’s divorced, then she’s been married’ (p. 116). Such conditionals are not purely epistemic because there is a(nother) relevant consequence relation underlying the truth of the conditional.

⁴ This example is due to Alistair Norcross in conversation.

has orbited the earth once or that if lightning strikes the church steeple, then a witch is to blame. Conditionals such as (7), if they are judged to be true, must be so judged for reasons that do not allow the conditionals of the geocentricist or the witch-hunter to be true as well. So, it should not be the case that it is the utterer's own credences that make (7) true. Note that treating (7) and other purely epistemic associations as true would not saddle the Consequence Account with all the problems of the purely epistemic Suppositional/Probabilistic Accounts, because it would not, for example, saddle the Consequence Account with Gallimore's Problem. (Laws of nature could serve to account for the truth of the relevant conditionals, even if no person was aware of those laws of nature. See section II.H of this chapter for further discussion.) It would only saddle the Consequence Account with the other half of the problem: that of allowing in conditionals that seem false, such as that of the geocentricist above. For that reason, it is important that the Consequence Account have a well-motivated standard for allowing conditionals such as (7) to be true for reasons other than the person's having certain beliefs—or else to reject them as technically false, though useful.

When considering the Consequence Account and the problems associated with purely epistemic conditionals, one might be tempted to say that (7) in this case is utterly false and end the matter at that. After all, last week's discussion about the TV show "Parks and Recreation" is not a consequence of my interlocutor's (now) thinking of Wednesday; such a consequence would be anachronistic. However, this reading of the situation strikes me as somewhat obtuse. When I utter (7), I am not intending to communicate that my actual discussion last week occurred as a consequence of my interlocutor's present thoughts. I am, rather, taking a potential piece of information—that my interlocutor is thinking of a conversation taking place last Wednesday—and linking that piece of information to another piece of information I have about what I was conversing about last Wednesday. There is a consequence relation in the vicinity, and it is between

that prior conversation and itself. The conversation on Wednesday *was* a conversation about “Parks and Recreation,” and so the following conditional is true:

(8) If you’re thinking of the conversation on Wednesday, you’re thinking of the conversation in which I was talking about the TV show “Parks and Recreation.”

This conditional is true, because (opaque contexts aside) thinking of A has thinking of the thing to which A is identical as a consequence. So, the case of (7) may be a case of a strictly false conditional, but its use can be explained as a proxy for a true conditional, (8), which (7) loosely and more briefly expresses. This result—that (7) is strictly false, but useful and used because of its relation to a true conditional—may seem to be a cost, but if so it is a minor one. We are able to explain the usefulness of conditionals such as (7) without being saddled with judging all purely epistemic conditionals to be true, because not all purely epistemic conditionals will be looser ways of expressing a definite description. For example, a superstitious sports fan who always happens to be wearing a certain pair of jeans when her favorite team loses may make the following purely epistemic association:

(9) If I wear these jeans, my team will lose their game.

Such a conditional is false according to the Consequence Account, because conditionals are not in general made true by relations between people’s beliefs. My believing that P has Q as a consequence is not sufficient for P’s having Q as a consequence. Contrary to the useful and loosely-correct case of (7) and (8) above, in the case of the purely epistemic (9) there is no consequence relation in the vicinity that makes a relevant conditional true, and so (9) comes out as flatly false, as it should.

Purely epistemic associations are not, in general, instances involving consequence relations in the sense at use in the Consequence Account, and so they are in general false according to this account.

5. Backwards evidential conditionals

There is a special case of what might be thought of as a purely epistemic conditional which, again, is common enough to deserve its own treatment. This special case is what I will call a

backwards evidential conditional: a conditional in which the antecedent is being presented as evidence for the consequent because the *antecedent* is a consequence of the *consequent*. Consider the case of (10), which is about one's neighbors, the Hernandez family:

(10) If the lights are on next door, then the Hernandez family is home.

The conditional (10) expresses a common use of the indicative conditional, using it to express a relation in which the antecedent is evidence for the consequent. Again, P's raising the subjective probability of Q is not enough for P to have Q as a consequence. So why does (10) seem to be a possibly true conditional? The lights being on certainly does not have as a *causal* consequence that the Hernandez family is home. In fact, if (10) is true, it seems to be because there is a causal relation in the other direction: the Hernandez family, when home, turns the lights on. It seems that (10) is a good use of the conditional just in case that causal relation holds—such that, for example, it is unacceptable or false in a scenario in which another family now lives in the home being referenced in (10).

As with purely epistemic associations in general, the Consequence Account judges (10) and other evidential conditionals lacking a consequence relation between the antecedent and consequent to be strictly false. Our use of the indicative conditional to express relations between our own beliefs is a secondary kind of use, parasitic on the standard case of a consequence relation. This judgment has the result that many acceptable-seeming evidential conditionals come out as strictly false. For example, a detective who is prone to discussing clues by means of indicative conditionals will end up making many a false statement. ('If the glove was in the parlor, then the gardener did it!' Etc.)

One might attempt to accommodate backwards evidential conditionals and other evidential conditionals that are helpful and good within the Consequence Account by using an analogy with logical entailment, but this approach is not advisable. When using a backwards evidential conditional, a person is expressing a kind of consequence relation between his or her beliefs. We

have already established that, in general, believing that the consequent follows from the antecedent is not enough to make a conditional true. However, we might draw an analogy here between the case of backwards evidential conditionals and logical consequence. In the case of logical consequence, Q follows from P as a matter of logic. Backwards evidential conditionals may operate according to something similar, where Q follows from P as a matter of reasoning. The idea is that what separates the true conditionals of this kind from the false ones is whether or not the reasoning is good, in a way that I will not try to specify here. Douven (2008) and Krzyżanowska et. al. (2014) present an account that treats all conditionals according to this measure.⁵ However, treating the ability for one to reason from P to Q as a genuine consequence relation is worrisome, because good reasoning with incomplete information can lead to false conclusions, and so modus ponens would be seriously undermined with this treatment of evidential conditionals. Treating such reasoning as a genuine consequence relation is therefore ill advised.

Treating evidential conditionals as false in general may seem to be a substantial cost for the Consequence Account (and perhaps the *only* significant cost), but there are reasons for thinking that this result is not a cost at all. Linguist Eve Sweetser (1991) presents a taxonomy of conditionals according to which epistemic conditionals (of which backwards evidential conditionals are one kind) are different in kind from the types of conditionals I have all along treated as standard, and therefore are deserving of a separate treatment. While, all else being equal, a unified account of the phenomena is preferable to a disunified account, it may be that not all else is equal when it comes to evidential conditionals. I did not want to prejudice the case by assuming from the outset that evidential conditionals deserve a different treatment from the kinds of conditionals dealing with

⁵ Here I am writing only about the (un)advisability of treating relations of good reasoning as one kind of consequence relation among many, but it is worth mentioning some of the problems of treating these relations as the sole measure of the acceptability/truth of a conditional, as Douven (2008) and Krzyżanowska et. al. (2014) do. The main problem with such an account is that it undesirably separates the truth/acceptability conditions of conditionals from the extra-mental world. These accounts suffer from Gallimore's Problem (for discussion of which, see II.H).

matters of causal, logical, nomological, or other dependence, but Sweetser's assessment is worth pointing out as a means of mitigating the apparent cost of treating evidential conditionals as false under the Consequence Account. It may even be that the considerations I give in favor of a worldly account themselves show the need for a separate account to treat evidential conditionals. Perhaps one should not expect an account to cover both evidential conditionals and the other, more worldly kinds of conditionals that the Consequence Account handles so well. In fact, this result may be good news both for the Consequence Account and for the Suppositional/Probabilistic Accounts that face so many worldliness-related problems. It may be that one can give an adequate account of evidential conditionals along the lines of, say, Douven's (2008) account, which seems to be a non-worldly counterpart of the Consequence Account, and together with the Consequence Account cover these two major kinds of indicative conditionals. In any case, it is clear that the Consequence Account does not entirely account for the good-seemingness of good-seeming evidential conditionals; it is less clear that this result is a cost for the Consequence Account, because a separate account may be necessary anyway.

6. Accidental universal generalizations

Sixth and finally, we have the case of conditionals that express accidental universal generalizations. Consider (11):

(11) If x is a living monarch, then x is over five feet tall.

Let us assume for our purposes that all living monarchs are over five feet tall and that this fact is not the result of some rule or other reason making it such that no person less than five feet tall could serve as monarch. It merely happens to be the case that there is a true universal generalization that all living monarchs are over five feet tall. (We are discussing accidental universal generalizations in

particular because they lack some reason why all x's have some feature, to rule out cases in which some other consequence relation accounts for the truth of the conditional.)

At an initial glance, (11) seems to be an acceptable, true conditional. In fact, people who have studied logic might be inclined to treat (11) as equivalent to the universal generalization. What does the Consequence Account have to say about it? Is there a consequence relation that underlies the truth of conditionals like (11)? Indeed there is, and the consequence relation is between the parts of the universal generalization. Being a living monarch has as a consequence that one is over five feet tall—not causally, of course, but because of the truth that all living monarchs are over five feet tall. This kind of consequence is similar to a case of a conditional involving a description of a single entity. Consider (12):

(12) If the person in the hallway is Michael Tooley, then the person in the hallway is a philosopher.

The conditional (12) is true because of the features of Michael Tooley. Because Michael Tooley *is* a philosopher, any identification of x with Michael Tooley is an identification of x with a philosopher. X's being Michael Tooley has x's being a philosopher as a consequence, not because Michael Tooley must be a philosopher, but because Michael Tooley is one. This same principle applies to accidental universal generalizations. Even though it is not the case that any living monarch must be over five feet tall, it is the case (within our supposition) that every living monarch is over five feet tall; and this accidental universal generalization generates a contingent consequence relation between being a living monarch and being over five feet tall that is similar to that between being a certain entity and being an entity with that thing's particular properties. Thus, conditionals truly expressing accidental universal generalizations are true according to the Consequence Account.

B. Valid and invalid inferences

Different accounts of indicative conditionals differ with respect to which inferences they judge to be valid. In this section, I discuss which inferences are deemed valid and which inferences are deemed invalid by the Consequence Account, as well as those that are invalid but useful for practical purposes. The conception of validity at play here is that an inference from A to B is valid just in case it is not possible for A to be true while B is false. Multiple sentences may be included within A.

1. *Modus ponens and modus tollens*

Modus ponens—the inference from A and ‘If A, then B’ to B—is central to our use of conditionals, and it is valid according to the Consequence Account. The reason modus ponens is valid is that, for any true conditional, there is a consequence relation between what is described in the antecedent and what is described in the consequent. Thus, if the antecedent is true (what is described in the antecedent obtains), then what is described in the consequent should also obtain, making the consequent true.

There are famous purported counterexamples to modus ponens, which seek to undermine its validity. One example is from Vann McGee (1985), concerning an election in which Anderson and Reagan are the only Republican candidates and Anderson is a distant third to the favored Reagan, who in fact wins the election:

If a Republican wins the election, then if it’s not Reagan who wins it will be
Anderson.
A Republican will win the election.
If it’s not Reagan who wins, it will be Anderson. (462)

The premises are true in the described scenario, but the conclusion seems false. It seems that, if Reagan loses, it will be someone else besides Anderson who wins. I diagnose cases such as McGee’s as involving a subtle shift in context. The reason we reject the conclusion is that we are considering

it full stop, rather than considering it in light of the supposition that a Republican wins the election. Shifting the context from the premises to the conclusion of an argument (or the input of the inference to its result) is a kind of equivocation, and thus purported counterexamples involving a shifting context are not genuine counterexamples. Lycan (2001) points out that dissolving the purported counterexamples by attributing a shift in context has the result that not all cases that superficially seem to be cases of modus ponens are genuine cases (see pp. 68-69), but that is a cost it seems necessary to bear in light of the purported counterexamples, and indeed the existence of equivocation in general entails this result.

Modus tollens—the inference from not-B and ‘If A, then B’ to not-A—is likewise a valid inference. If there is a consequence relation between A and B, then B’s being false is sufficient for A’s being false. Even though some consequence relations obtain contingently (such as the causal relation that obtains between rolling the die and the die coming up six), any relation where P has Q as a consequence is such that P is sufficient for Q. Thus, the falsity of the consequent in an indicative conditional is sufficient for the falsity of the antecedent; for if the antecedent were true, then the consequent would also be true. Since modus ponens and modus tollens are widely regarded as inferences worth preserving—if not inferences that are central to our use of conditionals—this result is a good one.

2. *Antecedent strengthening*

Antecedent strengthening—the inference from ‘If A, then B’ to ‘If A and C, then B’—is invalid according to the Consequence Account, and it is not a good inference for practical purposes. Consider again the causal conditional (4):

(4) If the vase is knocked off the counter, it will break.

Suppose that (4) is true of some vase. It does not follow that (13) is true:

(13) If the vase is knocked off the counter *and I catch it midair*, it will break.

Because the antecedent of the conditional does not always fully specify the situation at hand, it is not the case that the addition of any clause whatsoever to the antecedent will preserve the truth of the conditional. Antecedent strengthening is invalid, and this result is consistent with that of most other accounts besides the Material Implication Account.

3. *Transitivity*

Transitivity—the inference from ‘If A, then B’ and ‘If B, then C’ to ‘If A, then C’—is invalid according to the Consequence Account, though it may be good to use for practical purposes. The reason transitivity is invalid is that there is no guarantee that, if there is a consequence relation between A and B and a consequence relation between B and C, and there will also be a consequence relation between A and C. Cases in which the consequence relations between A and B and between B and C are of different types provide good counterexamples. For example, consider the case of (14)-(16):

- (14) If I move my knight to L6, April will move her queen to L6.
- (15) If April moves her queen to L6, April will have me in checkmate.
- (16) So, if I move my knight to L6, April will have me in checkmate.

The conditional (14) is a case of causal consequence. My doing one action will lead to April’s doing another action. Conditional (15), on the other hand, synchronously describes what state the board will be in, according to the rules, if April’s queen is at L6—namely, a state of checkmate. However, contrary to (16), my moving my knight to L6 does not, as a matter of the rules, put April in a position of having me in checkmate. Transitivity here fails, and so transitivity is not a valid inference. Nonetheless, many instances of transitivity will not involve a shift in the kind of consequence relation involved, and these might end up being good inferences. Some kinds of consequence

relation are themselves transitive (e.g., causation and logical entailment), and transitivity inferences involving only one of these kinds of consequence relation will be good inferences.

As Edgington (2014) notes, the fate of transitivity seems tied to the fate of antecedent strengthening, because antecedent strengthening is a special case of transitivity, as follows: If A and B, then A [tautology]; if A, then C; therefore, if A and B, then C. Thus, we should not be surprised that both inferences turn out to be invalid. Since most accounts besides the Material Implication Account treat antecedent strengthening as invalid, most should likewise treat transitivity as invalid.

II. What the Consequence Account does right

The Consequence Account succeeds in areas in which other major accounts fail. In this section, I discuss the judgments that the Consequence Account gets right, often in contrast with the major accounts.

A. The Consequence Account has intuitive appeal.

The Consequence Account has an intuitive appeal and a sort of naturalness shared only with Probabilistic/Suppositional Accounts. The claim that the conditional ‘If you drop the vase, it will shatter’ is true just in case dropping will have shattering as a consequence is very straightforward and believable. This conception of conditionals thus has a great deal of intuitive appeal. In contrast, the Material Implication Account is initially implausible and must be supplemented with a story about rules of assertion to diminish the initial bizarreness of some of its judgments. Similarly, the Possible Worlds Account involves technical machinery that keeps that view from being as intuitive as the Consequence Account.

B. The Consequence Account assigns truth values to conditionals.

In contrast with Suppositional/Probabilistic Accounts, the Consequence Account assigns truth values to conditionals. This feature counts in favor of the Consequence Account, because it fits well with the way speakers seem to use conditionals—to assert, disagree, etc. As discussed in Chapter Four: Truth Values, there are many initial reasons for thinking that indicative conditionals have truth values, and there are good reasons for rejecting various arguments to the contrary. Thus, while it has not been proven in this dissertation that indicative conditionals have truth values, the fact that an account assigns truth values to indicative conditionals counts in its favor as a way in which the account fits well with popular uses of conditionals.

C. There is no need for indeterminate truth values for conditional predictions whose antecedents do not obtain.

One of the benefits of the Consequence Account is how it treats conditional predictions whose antecedents do not obtain. Consider a case of a random number generator, set to generate an integer at random between 1 and 10.

(17) If I run the generator this afternoon, it will generate a six.

Suppose I do not run the generator this afternoon. Is (17) true or false, or... what? Some would say this conditional is undefined or that it lacks a truth value (for example, McDermott (1996)). Since I do not run the generator this afternoon, nothing determines the truth value of the consequent.

There is a certain intuitive appeal to the claim that the truth value of (17) is undefined in cases in which I do not run the generator this afternoon. Yet this claim would introduce some complexity to the logic of conditionals,⁶ and more importantly it seems false for certain conditionals.

⁶ But, see McDermott (1996), p. 5, for a partial working out of some of the implications for embedding conditionals with the possibility of undefined truth values.

As Brian Weatherson (2001) has pointed out, “Many theorists hold that indicative conditionals, especially those with false antecedents, lack truth-values. This cannot be right in general, since it is a platitude that $p \rightarrow p$ is true for every p ” (p. 212, with ‘ \rightarrow ’ as shorthand for the indicative conditional).

In order to accommodate Weatherson’s observation and retain T/F truth values for conditionals whose antecedents do not obtain, one could disagree with the claim that the truth value of (17) is undefined, call it false instead, and explain its falsity in terms of the lack of an appropriate consequence relation. In cases in which the antecedent is false, the contingent events of the world cannot fill in the appropriate consequence relation, and so the conditional will be true only if some other consequence relation makes it true.

For cases of true conditional predictions, the world fills in the answer. To use the familiar die-rolling case from Chapter Three, the rolling of the die *does* have six as a consequence, even though it might have led to any of five other numbers. However, in the case of (17), the world cannot fill in the actual causal story that happens to take place. There is no actual consequence that running the generator this afternoon has, much less a consequence of six, and so the conditional comes out false. No consequence relation exists between the antecedent and consequent of (17). However, this result does not mean that every future-tense indicative conditional with a false antecedent is false, for sometimes there will be another consequence relation between the antecedent and consequent—usually, some kind of strict dependence. For (18) below is uncontroversially true:

(18) If I run the generator this afternoon, I will run some bit of machinery this afternoon.

The world does not provide a consequence relation in the form of an actual event that has some other actual event as its causal consequence, but there is another consequence relation: one of logical consequence, which renders (18) true. This seems to be the right judgment.

One potential objection to this analysis is that it seems to entail that conditionals such as (17) mean something slightly different, depending on whether or not their antecedents are true. If the antecedent is true, we see what the world says in evaluating the conditional. If the antecedent is false, then some other kind of consequence takes precedence. Yet neither the meaning of a conditional nor its truth conditions should vary by whether or not their antecedents are true, and so this analysis (if the charge is correct) is unacceptable.

I agree that the meaning of a conditional should not depend on the truth value of its parts. Neither should the truth conditions. But neither of those things takes place in this account. No matter what the truth value of its antecedent, (17) means that my running the generator this afternoon has (will have) six as a consequence. No matter what the truth value of its antecedent, (17) is true just in case there is some consequence relation such that my running the generator this afternoon has (will have) six as a consequence.

No matter whether or not I run the generator, there are various options as to what the potential truth-maker of (17) could be. If it so happens that I do run the generator, then there will be some actual causal relation in the world that, if events turn out sixwardly, can serve as the truth-maker for (17). If, on the other hand, I do not run the generator, that does not alter the meaning of the conditional. It merely rules out a certain kind of potential truth-maker: an actual causal relation between two events in the world, one of which (the cause) is a generator-running event and the other of which (the effect) is the event of getting six. The meaning of the conditional is constant, though the availability of certain truth-makers will cause its truth value to be different in different contexts. But the varying availability of truth-makers is perfectly normal. Someone who kills the very last dog has not altered the meaning of 'dogs exist'; such a person has merely changed its truth value (in addition to doing something that is no doubt morally significant). Similarly, my decision not to

run the generator this afternoon does not alter the meaning of (17); I have merely eliminated the possibility of my making (17) true. Such a result should not be a cause for concern.

D. The Consequence Account explains the ambivalence of certain cases.

Another benefit to the Consequence Account is that it can help explain one's ambivalence about the truth value of conditionals in certain tricky cases. Mackie (1973), pp. 107-108, presents a case of a father and child at a zoo. The father says to the child,

(19) If you poke your finger into that monkey's cage, you will get it nipped off.

We are then asked to evaluate the truth of that conditional relative to different possible outcomes.

One possible outcome is that the child sticks a finger in the monkey's cage, and the monkey remains placid; but, after the child pulls the finger back out of the monkey's cage, a bird swoops down and nips it off. The antecedent and the consequent of the conditional in this case are both true. Mackie thinks we should for that reason say the conditional is true. (I have argued for the invalidity of that reasoning in Chapter Three; some TT conditionals are false.)

Mackie also claims that a connection-type account should judge this conditional to be false. Yet there is some unease in evaluating (19), and I think that, contrary to Mackie's claim that such an account should judge the conditional false, the Consequence Account lends some support to either claim (truth or falsity)—and perhaps, in the end, the most support to ambivalence. The consequent of (19) is true, but not for the reason the father suggested it would be, and so it is unclear whether his conditional prediction should count as a success or not. Such is the character of my unease in evaluating (19). I submit that the ambivalence with respect to the conditional in this case is borne out by the Consequence Account. The nipping of the finger is a later part of the causal chain that includes (as a near prior event) the child's sticking the finger in the cage. But the sticking the finger in the cage is not a *direct* cause (which I mean in an intuitive sense) of the finger-nipping. So it is

unclear whether or not it should count as a consequence of the fulfillment of the antecedent. The problem is not merely that the father's reasons for being confident of (19) were mistaken, but that the conditional itself is meant to express two (in this case) events that relate to each other in some way, and it is not clear that these two events relate to each other in the right way. Our theory predicts the unease that is a part of the felt response to this conditional. Such a phenomenon, then, counts in favor of the Consequence Account over other accounts that do not predict or explain this unease.

E. The Consequence Account makes the right judgments about TT conditionals.

In Chapter Three: TT Conditionals, we found that the major accounts make the wrong judgments about two phenomena surrounding TT conditionals (conditionals whose antecedents and consequents are both true):

(P1) There are some false TT conditionals. (There are some unacceptable BB conditionals.)

(P2) There are some cases in which the fact that the antecedent and consequent come out true *does* seem to make the whole conditional true, even when the consequent's truth is not entailed or necessitated by the antecedent.

The Consequence Account judges as true just those TT conditionals for which there is an appropriate consequence relation. So not all TT conditionals come out true, but the intuitively true ones do.

Consider the following conditionals from Chapter Three, renumbered here:

(20) If grass is green, then Obama is President.

(21) If I exist, then the Seahawks won the 2014 Super Bowl.

(22) If Obama is President in 2014, then the star Betelgeuse will someday go supernova.

In cases such as the false TT conditionals (20)-(22), there is no consequence relation between the antecedent and consequent. Obama's being President is not a consequence of grass's being green,

etc. The Consequence Account has an explanation for the fact that these conditionals, despite our belief in the truth of their component parts, seem defective and false—namely, that they *are* false. The Consequence Account judges them to be false because there is no relation such that their consequents are the consequences of their antecedents. Speaking in terms of probabilities, while it may seem odd to say that it is highly unlikely that if grass is green, then Obama is President, it *is* unlikely that there is a relation by virtue of which Obama's being President is a consequence of the fact that grass is green. And this unlikelihood fits well with the judgment that (20) and the other conditionals with irrelevant components are false.

The Consequence Account also deals adequately with the false TT conditionals (23) and (24):

- (23) If it is windy, then I will hear the wind chimes chiming.
- (24) If Leslie is bitten by a tick, she will contract Lyme disease.

In each of the scenarios described for these conditionals in Chapter Three, there is no consequence relation between the antecedent and the consequent of the conditional. With respect to (23), there is no causal or other consequence relation between its being windy and my hearing the wind chimes; I hear the chimes for another reason. With respect to (24), there is no causal or other consequence relation between Leslie's being bitten by a tick and her contracting Lyme disease. So, (23) and (24) come out false according to the Consequence Account. Thus, the Consequence Account adequately accommodates (P1).

The Consequence Account likewise accounts for (P2). In cases of successful conditional predictions, such a consequence relation does exist. Consider again the TT conditional (25), renumbered here:

- (25) If Sophia rolls the die, the die will come up six.

The consequence relation in this case is a causal one. Though it is possible that Sophia could roll any number one through six, in fact her rolling the die causally results in her rolling a six. Her rolling the

die has her rolling a six as a consequence, and so the conditional comes out true according to the Consequence Account. The Consequence Account correctly predicts the truth value of this conditional and, by generalization, conditionals of this kind. If, however, the events in the TT conditional are not causally or otherwise consequentially related, then the conditional is false, as we saw in the case of (23) and (24), as well as (5) above:

(5) If I roll a six, your cat will die tonight.

The Consequence Account thus has principled reasons for accepting the intuitively true TT conditional predictions while rejecting those such as (5) which seem false.

Others have described the phenomenon of (P2) by saying that, in cases in which the antecedent is true or what it describes obtains, the actual world “fills in” in the truth value of the conditional by providing the truth value of the consequent, which then determines the conditional’s truth value. Yet if the foregoing analysis is correct, this description is not entirely right. The world only “fills in” the truth value in cases in which there exists a contingent consequence relation (such as causation). Noticing this fact allows us to embrace an account according to which not all TT conditionals are true (which fits the phenomena of (20)-(24)), and yet all true conditional predictions are true. Thus, both (P1) and (P2) are accounted for. Not only that; the Consequence Account also explains the intuitive appeal of those theories according to which all TT conditionals are true. Most often, we use conditionals only in cases in which we do not know whether or not the antecedent will obtain—cases such as conditional predictions where the antecedent and the consequent are part of the same causal chain, for which all TT instances are true. Given the ubiquity of cases of this kind of conditional prediction, it is understandable that theorists would be tempted to draw the conclusion that all TT conditionals are true. Since the Conditional Account illuminates the fact that all TT conditional predictions whose components are a part of a causal chain are true, while correctly predicting that not all TT conditionals whatsoever are true, it explains in an intuitive way

the common mistake of denying (P1); this mistake stems from the ubiquity of the kind of conditional all of whose TT instances are true.

F. The Consequence Account gets the Sly Pete case right.

Another feature of the Consequence Account that counts in its favor is its response to Gibbard's Sly Pete Case (discussed in Chapter Four, section IV). Recall that Jack and Zack believe contradictory conditionals in the Sly Pete Case, labeled with the initial of the person who believes each conditional:

- (J) If Pete called, he lost.
- (Z) If Pete called, he won.

If both these conditionals are acceptable in the same context, then one must either treat the truth value of a conditional as radically indexed to the epistemic state of the person believing the conditional or else do away with truth values altogether. The Consequence Account gives us a principled reason for rejecting Gibbard's assessment that both (J) and (Z) are acceptable, thus avoiding Gibbard's dilemma. According to the Consequence Account, (J) is true in this case, and (Z) is false. (J) is true because there is a consequence relation generated by the rules of the game, given the hands of the two players, such that Pete's calling has his losing as a consequence. (Z) is false in this case, because there is no consequence relation between his calling and his winning. If Pete called, he lost, based on the rules of the game; and in no way is Pete's winning a consequence of his calling. Because the Consequence Account can avoid the problematic dilemma that is the intended result of Gibbard's case, the Consequence Account adequately deals with Gibbard's Sly Pete Case. Not only does the Consequence Account respond adequately to the case; it also solves the case in a way that fits well with initial responses to (J) and (Z). It seems intuitive to say that (J) is true, and (Z) is false, and the Consequence Account explains why this response is not only intuitive but correct.

Thus, the ability of the Consequence Account to respond to Gibbard's Sly Pete Case counts as a point in its favor—not only avoiding the dilemma but also explaining initial responses to the case.

G. The Consequence Account deals with Barnett's Conscious Being Case.

Recall David Barnett's (2012) Conscious Being Case from Chapter Four: In Favor of Truth Values. Similarly to Gibbard, Barnett presents a situation in which we are meant to judge that two people who hold contradictory conditionals in the same scenario are both justifiably certain of those conditionals. Barnett says it is reasonable for him to accept (26) with full certainty:

(26) If there is only one conscious being, David Barnett is it.

However, those of us who are not David Barnett must be certain of the contrary of (26):

(27) If there is only one conscious being, David Barnett is not it.

My certainty in (27) is based on my belief in my distinctness from Barnett as well as my belief in the following:

(28) If there is only one conscious being, I am it.

Since I am certain that I am a conscious being, (27) is eminently reasonable to me, just as (26) is to David Barnett, and one would not fault either of us for accepting our respective conditionals.

Moreover, on a certain conception of justifiable belief, if one is justifiably certain of P, then P is true.

Yet (26) and (27) are contradictory. If they are both true at the same time, then the Principle of Conditional Non-contradiction is violated. Barnett concludes that indicative conditionals lack truth values, though his case has not ruled out the result that we should treat conditionals as radically-indexical truth-valued sentences.

I discussed a general approach toward responding to Barnett's Conscious Being Case in Chapter Four, but here I will present a response supplemented by the resources available to the Consequence Account. If the Consequence Account is correct, it is facts about relations in the world

that make true conditionals true. And facts about a two-or-more-person world do not necessarily contain facts about who exists if only one of them does. According to the Consequence Account, the lack of an appropriate consequence relation results in a false conditional, and so it turns out that I am not justifiably certain of (27). My response to Barnett's Conscious Being Case is that the need to choose to accept either (26) or (27) is an illusion. They can both be false, and probably they both are, because in a multiple-person world such as ours, there is plausibly no consequence relation between the existence of only one person and the existence of any particular person. Barnett's (2012) example, Barack Obama and Hu Jintao disagreeing about (29) and (30), fares no better.

(29) If only one conscious being exists, it is Barack Obama.

(30) If only one conscious being exists, it is not Barack Obama.

The reason, again, is that nothing about the world, as far as we know, settles the facts about who exists, if only one of these leaders exists. Both sentences are false. Judging (26) and (27), as well as (29) and (30), to be false means denying the principle of Conditional Excluded Middle, according to which, for any A and B, either 'If A, then B' is true/acceptable, or else 'If A, then not-B' is true/acceptable. The Consequence Account rejects Conditional Excluded Middle, because sometimes both a conditional and its contradictory are false, as in the present cases—because sometimes neither a statement nor its contradictory is a consequence of the antecedent. Thus, the Consequence Account gives a principled reason for rejecting Conditional Excluded Middle and for rejecting the conditionals that Barnett uses in his case. Since all of them are plausibly false/unacceptable, the problem they produce never arises.

Such a result is, admittedly, counterintuitive. It seems that I should be justifiably certain that, if only one conscious being exists, I am it. But if our conditionals are about the world, then their truth values should depend on the way the world is. And facts about the world do not settle what follows from every antecedent. The claim that I am justified in accepting, or even being certain of, (27) is illusory. (26) fares better than (27) only insofar as I have little reason to believe that, in a

world of billions, the question is settled in Barnett's favor. One way to mitigate the counterintuitiveness of this result would be to embrace an alternative treatment of evidential/epistemic conditionals, as discussed with reference to Sweetser (1991) above (section I.A.5). Plausibly (26) and (27), as well as Zack's conditional (Z) from Gibbard's Sly Pete Case, express evidential conditionals, and if evidential conditionals deserve a separate account from worldly ones, these conditionals could end up being acceptable within such an account.

H. The Consequence Account does not suffer from Gallimore's Problem.

Another advantage that the Consequent Account has over alternative accounts is that it does not suffer from Gallimore's Problem, as Suppositional/Probabilistic Accounts do. (See Chapter Four, section VI, for further discussion of Gallimore's Problem.) Gallimore's Problem is that, if an account does not closely tie truth or acceptability of conditionals to relations in the world, then there will be intuitively true/acceptable conditionals which that account judges to be false/unacceptable. For example, take any conditional made true by laws of nature that are unknown to the person considering the conditional, as (31) would be to anyone who knew nothing of helium or balloons:

(31) If I let go of a helium-filled balloon, it will rise.

The conditional (31) would be unacceptable to anyone ignorant of the properties of helium-filled balloons or the relevant laws of nature—or, more perspicuously, anyone whose subjective credences were such that it was not the case that the probability of a helium-filled balloon's rising, given that one lets it go, was high. Yet (31) is intuitively acceptable and true, and it is a major strike against an account if it does not treat (31) as acceptable or true. ('Acceptable' here means acceptable in general—that is, good for rational acceptance—rather than acceptable with respect to someone's epistemic state.) The Consequence Account does not suffer from Gallimore's Problem, because the relevant laws of nature are what make (31) and conditionals like it true, regardless of anyone's

credences. According to the Consequence Account, conditionals such as (31) are made true by consequence relations in the world, and so this account entirely avoids Gallimore's Problem.

I. The Consequence Account explains certain peripheral uses of conditionals.

As I discussed briefly in the introduction to this dissertation, there are certain kinds of conditionals that are considered by many to be non-standard. A non-standard conditional is one that does not straightforwardly follow the acceptability or truth conditions that govern the majority of conditionals. While a good account need not—and should not—treat these kinds of conditionals as standard, it nonetheless should account for our use of them. If there are uses of 'if' that do not straightforwardly fall under the proposed account, then some explanation must be given of why we use 'if' in that way, preferably in light of its standard use. The Consequence Account can accommodate two prominent kinds of non-standard conditionals in this way.

One kind of non-standard conditional is a biscuit conditional (or more generally, a speech act conditional), exemplified by Austin's (1961) 'There are biscuits on the sideboard if you want some'. A biscuit conditional appears to be non-standard, because the location of the biscuits is neither presented nor understood as being conditional on the hearer's desire for some biscuits. Rather, the biscuit conditional extends an offer of biscuits to the hearer. Furthermore, the non-'if' clause of the biscuit conditional (it does not seem quite right to refer to this clause as the consequent) seems to be asserted when one asserts a biscuit conditional, while the consequent of a standard conditional is not asserted when one asserts a standard conditional. These reasons seem sufficient for treating biscuit conditionals and other speech act conditionals as non-standard. What, then is going on when one uses a biscuit conditional? In a biscuit conditional, the 'if'-clause is used to say under what circumstances the other clause is relevant to the hearer. If there are biscuits on the sideboard, then there are biscuits on the sideboard whether or not the hearer wants them; but the

information that there are biscuits on the sideboard is offered under the condition that the hearer wants biscuits. A biscuit conditional thus involves a sort of hedging, not on the offer of the biscuits (or the information about the biscuits), but on the implication that this offer (or information) is desired by the hearer. So, how does the hedging of the implication that the offer (or information) is desired by the hearer relate to the standard use of the conditional as expressing a consequence relation between the antecedent and the consequent? I submit that the biscuit conditional can be seen as elliptical for an expanded version of the conditional, which does involve a consequence relation. One candidate for an expanded version is (32):

(32) If you want some biscuits, then you will want to know this: there are biscuits on the sideboard.

The expanded version of the biscuit conditional is literally true according to the Consequence Account. A person's wanting some biscuits does plausibly have as a consequence that that person wants to know about the availability of biscuits. Furthermore, (32) seems to accomplish what the original biscuit conditional was meant to accomplish. (32), however, is not terribly brief, and so 'you will want to know' is dropped for brevity. Thus, while the biscuit conditional, non-expanded, does not turn out to be literally true, its use is explained in terms of its relation to this expanded version, which is a standard use of the conditional. Such a story seems consistent with our use of biscuit conditionals, and it explains this use in terms of the Consequence Account's description of standard conditionals.

Another kind of non-standard conditional is semifactuals, which some call non-interference conditionals: conditionals including 'still' or 'even', or whose plausibility would benefit from the addition of those words. Douven (2008) and Burgess (2004) are examples of philosophers who deem these kinds of conditionals as non-standard and so do not seek to accommodate them in their accounts of indicative conditionals. Examples are as follows:

(33) Even if I am ill, I will (still) attend the ceremony.

(34) Even if you complete all assignments, you could (still) fail the class.

The main reason for treating semifactuals as non-standard conditionals is that asserting a semifactual seems to involve asserting its consequent. Asserting a standard conditional, on the other hand, does not seem to involve asserting its consequent. The Consequence Account can explain the use of ‘if’ in semifactuals in light of the use of ‘if’ in standard conditionals in the following way: semifactuals gain their force by means of flouting the expectation of a consequence relation. ‘Even if P, Q’ expresses that P does *not* have *not-Q* as a consequence, contrary to what one might expect. The semifactual is used to deny the existence of a certain expected consequence relation. In our examples above, (33) is used to deny that my being ill will lead to my missing the ceremony, and (34) is used to communicate that completing all assignments is not sufficient for passing the class. Thus, the Consequence Account accommodates the use of semifactuals and biscuit conditionals—two kinds of non-standard conditional—in terms of its account of the standard conditional, which is a point in its favor.

J. The Consequence Account respects relevance.

Unlike all the other major accounts, the Consequence Account respects the pre-theoretical judgment that the antecedent and consequent of a conditional should be relevant to each other in a true conditional, because two things being linked by a consequence relation is a means of their being relevant to each other. Even if a consequence relation between P and Q is artificially imposed by the rules of a game, those rules render P and Q relevant to each other in their specified way. Other major accounts do not respect relevance between the antecedent and the consequent. The Material Implication Account, which operates on truth values of the component parts, does not require relevance between the antecedent and the consequent. Neither do Suppositional/Probabilistic Accounts, since two things need not be relevant to each other in order for their probabilities to be

related in the right way. The Possible Worlds Account likewise fails to include relevance, because two things need not be relevant to each other for the one to be true at the nearest possible world(s) in which the other is true. The Consequence Account alone accounts for the desire for relevance by locating relevance in the meaning and truth conditions of conditionals, rather than explaining the desire for relevance in terms of assertability.

K. The Consequence Account does not have to tell a separate story about assertability.

Unlike the Material Implication Account, the Consequence Account does not have to tell a story about why the assertability of conditionals differ from their acceptability, because the two do not differ. Nothing besides the likelihood of being true is required for a conditional to be assertable in general, if its truth conditions are what the Consequence Account deems them to be.

III. Objections and Replies

In this section, I address potential objections to the Consequence Account or to connection-type accounts in general from H.P. Grice, J.L. Mackie, and Dorothy Edgington.

A. Grice's objection to connection-type accounts

Grice (1989) is concerned with defending the Material Implication Account against competing accounts that require that there be some kind of connection or relation between P and Q besides (or instead of) the truth-functional one. The Consequence Account posits just such a connection, and so Grice's argument can be seen as an argument against the Consequence Account. Grice argues against the claim that a conditional means more than the material conditional by arguing against the claim that the indicative conditional can never be asserted just on truth-functional grounds. If something more than 'not-P or Q' were meant by the conditional, then the

mere belief that not-P or Q could never be grounds for asserting 'If P, Q'. Grice presents some cases and argues that, in these cases, the mere belief that not-P or Q is good reason for asserting 'If P, Q'.

Grice presents two kinds of case to support his claim that there are sometimes good truth-functional reasons for asserting 'If P, Q' in the absence of a further connection. (See Grice (1989), pp. 59-60.) Grice's first is a case of declaring 'If P, Q' on the basis of the knowledge that P and Q. He asks us to imagine a speaker who knows that Smith is working in the library and says,

(34) I know just where Smith is and what he is doing, but all I will tell you is that if he is in the library he is working.

Since, as I have argued in Chapter Three, not all TT conditionals are true, the knowledge that the antecedent and consequent are true is not sufficient grounds for asserting the conditional. This conditional is not a case of a conditional that is assertable just on truth-functional grounds, because it is not a case of a conditional that is assertable. Pace Grice, the asserted conditional *does* seem to me to carry the implication (or entailment) that there is some connection between Smith's being in the library and his working. Even if I know that the speaker knows Smith's location and what he's doing, I can imagine the speaker uttering Grice's conditional just in case the speaker has the further belief that Smith only ever goes to the library to work, and it is harder to hear the conditional as assertable in cases lacking this connection. The Gricean will treat the conditional in (34) as misleading-but-assertable, but it strikes my ear as bizarre, absent any consequence relation between Smith's being in the library and his working.

The second kind of case Grice presents has to do with uttering conditionals within the context of a logic puzzle or a game. Grice contends that, within the rules of a logic puzzle, sentences such as 'If Jones has black (pieces) then Miller has black pieces too' are assertable on mere truth-functional grounds. His other case is a case of uttering a conditional within a game of bridge, and the Consequence Account's response to both is the same: the rules of games generate consequence

relations, just as in the case of (3) ‘If Leon moves his queen, he will check-mate’. These conditionals come out as true according to the Consequence Account, and not for truth-functional reasons. The truth of such conditionals depends on contingent features of the game in question, and the Consequence Account accurately predicts this fact. The consequent is, for lack of a better term, a rule-generated consequence of the antecedent. The game places certain constraints on reasoning, so in another sense the consequence is a logical consequence, within the artificially-imposed logic of the game. So, Grice’s charge that these are cases in which the speaker has mere truth-functional grounds for asserting a conditional turns out to be misleading. The reason the truth-functional grounds are grounds for assertion in this case, as they would not be in others, is that within the context of the game there is a relevant consequence relation that connects the truth values of the component parts.

Mackie (1973) points out that connection-type accounts must have the ability to accommodate logic puzzles and other games, though he writes of this feature in a somewhat negative light. He says,

If his account is to cover all acceptable if-sentences, [the proponent of the a connection-type account] must take ‘being a consequence’ in a wide enough sense to include such artificially constructed cognitive consequences [as in the trivia game]; and he must not dismiss as absurd any if-sentence, however bizarre, however apparently unrelated its antecedent and consequent may be, but content himself with saying that it is *prima facie* unacceptable, that it doesn’t make sense as an ordinary if-sentence until some consequential connection is filled in. (68)

As Mackie notes, in cases in which there really is some appropriate connection between the antecedent and the consequent, an account of the indicative conditional should judge such a conditional to be true or acceptable—even if the connection is unexpected or artificially imposed by the rules of a game. Such a result does not, I submit, leave the proponent of the Consequence Account in an untenable situation with respect to the credence she puts in various “bizarre,” as Mackie puts it, conditionals. The proponent of the Consequence Account should put just as much

credence in the conditional ‘If P, Q’ as she does in the claim that there is some relation by virtue of which Q is a consequence of P. In cases in which it is prima facie implausible that there is such a consequence relation between P and Q, the proponent of the Consequence Account should give negligible credence to the conditional. Far from being a cause for concern, this result is what one should want and expect from an account of conditionals.

B. Mackie’s argument that connection-type accounts are circular

Mackie (1973) is another of the few who critique connection-type accounts in any direct way (as opposed, for example, to arguing against them in effect by arguing that indicative conditionals lack truth values). Two of his criticisms merely do not stick to the Consequence Account, and the third I will discuss at greater length. First, Mackie objects that it is implausible that the kind of ‘if’ present in an indicative conditional changes with the kind of consequence involved. “It would be almost as absurd,” he says, “to suggest that what conjoins two legal provisions must be a legal ‘and’” (pp. 82-83). I agree completely, and my version of the Consequence Account does not have this feature; ‘if P, Q’ means the same no matter what kind of consequence makes (or could make) the conditional true. Every conditional ‘If P, Q’ means that P has Q as a consequence, and the meaning of ‘if’ does not change with the various truth-makers for the conditional (causal relations, logical entailment, etc.), and more than the meaning of ‘three’ changes when items of different categories add up to three (three events, three puppies, three ideas...). Different kinds of consequence relation make the conditional true in different cases, but they do not change the meaning of ‘if’. In every case, ‘If P, Q’ means that P has (had, will have) Q as a consequence.

A second objection from Mackie is that the Consequence Account does not accommodate certain peripheral uses of conditionals, such as biscuit conditionals or semifactuals. As I have

discussed above (section II.I), the Consequence Account does indeed accommodate these kinds of non-standard conditionals, and so this objection has already been answered.

The objection of Mackie's that deserves further discussion here is, third, the objection that the account is viciously circular. Says Mackie, "Can we explain what it is for P to *ensure* Q, or for Q to be a *consequence* of P, without using the very conditionals that we were hoping to analyse?" (p. 83, italics in original). The objection is that conditionals are used—or perhaps that they must be used—in illuminating the very concept of consequence, and so the explanation of conditionals in terms of consequence creates a vicious circle of explanation.

Vicious circularity would indeed be a problem, but it is not a problem the Consequence Account suffers from. Here we may take a cue from those who are working on the metaphysics of dependence. Various kinds of consequence relations are relations of dependence, such as metaphysical dependence or causal dependence. In fact, I had been tempted previously to call my account a kind of 'Dependence Account'. (Such a term fits ill with conditionals such as 'If P, then P' and logical consequence in general, which is sometimes symmetric.) As noted above, many contemporary philosophers take dependence as primitive. (See, again, Fine (2001), Schaffer (2009), Rosen (2010), and Sider (2014).) Such philosophers gesture at dependence with examples and seek to describe some of the logical features of dependence relations (in which 'if' no doubt appears), but dependence is seldom defined or explained in terms of its logical features. Though 'if' no doubt features in a productive discussion of these consequence relations—as one should expect if they underlie our ubiquitous 'if'—I submit that consequence does not need to be explained or defined in terms of the conditional. The notion of consequence is graspable on the basis of a short list of kinds of consequence relations, as I have given. Such a gesture may be sufficient for an explanation of consequence, as it is considered to be in cases of metaphysical dependence. But it is worth exploring whether or not the items on the list require 'if' in their explanation. Dependence relations are, let us

suppose, primitive, and logical consequence can be explained in modal truth-functional terms (for example, using ‘not’ and ‘and’) without reference to conditionals. The literature on causation likewise includes plenty of options that do not analyze causation in terms of any kind of conditional. (See Schaffer (2014).) These kinds of consequence seem to cover most of the central cases of indicative conditionals, and none clearly relies on the conditional in explaining the consequence relation.

Such a discussion may raise the further question, in virtue of what are these disparate kinds of relation all *consequence* relations? It will not, of course, do to say that what unites them is their aptness for generating conditionals. I think that here all that can be appealed to is the intuitiveness of the concept and stable, shared judgments about what does and does not belong on the list of consequence relations. Such an appeal may be disappointing, but I think it is all one should expect from such a basic concept as the concept that underlies conditional speech and reasoning. It may be that the concept of consequence cannot be explained in any illuminating non-circular way. However, given the paradox of analysis, we should not be surprised if circularity worries arise. In fact, I think we should be more surprised if it turned out that our oft-used conditional were best explained in terms of an esoteric apparatus such as possible worlds.

C. Edgington’s argument against connection-type accounts

Edgington (1986) presents, almost as an aside, an argument against what she calls “strong connection” accounts—that is, accounts according to which there must be some sort of (strong) connection between the antecedent and the consequent in order for the conditional to be true. While Edgington does not define or describe any strong connection accounts, I assume that my own Consequence Account is in the spirit of these accounts. Edgington’s objection is that a strong connection is lacking in certain cases of standard conditionals, for example, in 50-50 probability cases such as tossing a fair coin. She says,

Non-truth-functional accounts of the truth conditions of conditionals demand some sort of ‘strong connection’ between antecedent and consequent for the conditional to be true. Such a connection is clearly lacking in ‘If you toss this (fair) coin, it will land heads’. On such accounts, the conditional is then certainly false. It should have probability 0. But surely, if someone is told ‘the probability is 0 that if you toss it it will land heads’, he will think it is a double-tailed or otherwise peculiar coin. (19)

Edgington’s objection involves two elements, both of which deserve treatment. The first is the claim that a connection-type account cannot treat ‘If you toss this (fair) coin, it will land heads’ as anything but false, since a connection between the antecedent and the consequent is, according to Edgington, “clearly lacking.” It should be clear by now what response the Consequence Account makes to such a claim. According to the Consequence Account, the conditional in question is true just in case you toss the (fair) coin, and it does land heads. (The designation of the coin as fair rules out other potential consequence relations, because heads is not the only possible outcome.) Should someone toss the coin, tossing the coin will have heads or tails as a result, as a *consequence*, and so a heads outcome would make Edgington’s conditional true according to the Consequence Account. Prior to tossing the coin, and absent any special foresight, it would be a mere guess to assert the conditional in question, but the conditional is not necessarily false (or “certainly false”) according to the Consequence Account. Prior to tossing the coin, one should assign the conditional exactly the subjective credence one assigns to sentences that have an equal chance of turning out true or false.

The other element of Edgington’s objection is that stating the probability of the conditional is misleading. As discussed above, the conditional could turn out to be true or false, but suppose that the conditional, in fact, is known to be false. It seems wrong in that case to say that there is a probability of zero that, if you toss the (fair) coin, it will land heads. The reason this statement seems wrong is that it misleads one about the likelihood of a heads outcome to the coin toss. Yet this statement is not objectionable, once one understands that it expresses the known truth value of the conditional. It is only when one takes it to express the conditional probability of the consequent, given the antecedent, that the statement is misleading. There is no more problem here with the

expression of zero probability for a false statement whose prior probability was higher than zero than there is in general for the treatment of any such false statement as having zero probability.

IV. Conclusion

The prospect of a good connection-type account of indicative conditionals has been, for the most part, either ignored or objected to in recent decades. Yet the Consequence Account shows that one should not rule out the possibility of such an account. The Consequence Account succeeds where the major accounts face significant challenges. It fits well with and explains what speakers are up to when using conditionals, while respecting the well-documented intuition that there must be some sort of connection between antecedent and consequent. Given its short list of weaknesses and long list of strengths, the Consequence Account deserves strong consideration.

CONCLUSION

TRUTH AND CONSEQUENCE

Conditionals are at once familiar and obscure, easy to use and difficult to account for. In this dissertation, I have illuminated several aspects of indicative conditionals that the major accounts have gotten wrong. As I have argued, the Material Implication Account cannot make use of conversational implicature to explain the sense that there should be some connection between the antecedent and consequent of a conditional. Furthermore, the Material Implication Account fails an easy test for a theory of truth conditions. In addition to these problems, versions of all the major accounts fail to recognize that certain TT or BB conditionals are unacceptable. None of the major accounts adequately accommodates all the phenomena related to TT/BB conditionals.

Yet there is hope for an account, because it turns out that many have been wrong about the prospects for assigning truth values to indicative conditionals. Because of arguments by David Lewis, Dorothy Edgington, Alan Gibbard, and David Barnett, many have concluded that indicative conditionals must lack truth values, in spite of pre-theoretic intuitions to the contrary. These arguments can adequately be answered, vindicating the thought that indicative conditionals can be true or false. The arguments of Lewis and Edgington mistakenly rule out possible assignments of acceptability/unacceptability or truth/falsity to various conditionals, based on the truth values or probabilities of their component parts. Gibbard's argument mistakenly treats an unacceptable conditional as acceptable. Though Gibbard tells a story in which one might not fault two people for concluding two contrary conditionals, we can see that one of the two conditionals is false in Gibbard's scenario, dissolving his dilemma between giving up truth values and radically indexing

conditionals to their utterers' epistemic states. Barnett's argument mistakenly treats certain conditionals as acceptable, when there is no fact of the matter about whether or not their consequents are true, given their antecedents. His argument likewise fails to show that indicative conditionals lack truth values.

With the possibility of a truth-valued account comes the chance for an account that respects the well-documented sense that there should be some connection between a conditional's antecedent and consequent. One way to fill out the details of such an account is with the Consequence Account, according to which 'If A, B' is true just in case there is some relation by virtue of which B is a consequence of A. This account has the result that you and I can assert, deny, and disagree about conditionals, just as we are accustomed to doing. (This result should come as a particular relief to philosophers, whom some have thought to be primarily in the business of arguing about conditionals.) Our conditionals, which seem mostly to be about the world beyond our own minds, end up being made true (if at all) by features of the extra-mental world—by causes and effects, logical entailments, the rules and conventions that people have brought into existence, and many other kinds of relation. The Consequence Account explains the phenomena surrounding TT conditionals that the major accounts fail to accommodate, giving it a distinct advantage in terms of fit with how people use conditionals. Finally, the Consequence Account alone validates our desire for a strong connection between the antecedent and consequent, and it has many advantages besides these, as discussed in Chapter Five. The Consequence Account, a kind of account more or less ignored for decades, stands as a peer with the major accounts, and it deserves strong consideration. With the Consequence Account, it seems that we can have it all when it comes to an account of indicative conditionals: both truth and consequence. If that's not enough, I don't know what is.

BIBLIOGRAPHY

- Adams, E. (1965). A logic of conditionals. *Inquiry*, 8, 166-197.
- Adams, E. (1966). Probability and the logic of conditionals. In J. Hintikka & P. Suppes (Eds.), *Aspects of Inductive Logic* (pp. 256-316). Amsterdam: North Holland.
- Adams, E. (1970). Subjunctive and indicative conditionals. *Foundations of Language*, 6, 89-94.
- Adams, E. (1975). *The Logic of Conditionals*. Dordrecht: Reidel.
- Appiah, A. (1985). *Assertion and Conditionals*. Cambridge: Cambridge University Press.
- Austin, J.L. (1961). Ifs and cans. *Philosophical Papers*. Oxford.
- Bach, K. (1999). The myth of conventional implicature. *Linguistics and Philosophy*, 22(4), 327-366.
- Barker, S. (1997). Material implication and general indicative conditionals. *The Philosophical Quarterly*, 47(187), 195-211.
- Barnett, D. (2006). Zif is if. *Mind*, 115, 519-566.
- Barnett, D. (2012). Future conditionals and DeRose's thesis. *Mind*, 121, 407-442.
- Belnap, N. D. (1970). Conditional assertion and restricted quantification. *Nous*, 4(1), 1-12.
- Bennett, J. (1988). Farewell to the phlogiston theory of conditionals. *Mind*, 97, 509-527.
- Bennett, J. (1995). Classifying conditionals: The traditional way is right. *Mind*, 104, 331-344.
- Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford: Clarendon Press.
- Blackburn, S. (1986) How can we tell whether a commitment has a truth condition? In C. Travis (Ed.), *Meaning and Interpretation* (pp. 201-232). Basil Blackwell.
- Bradley, R. (2000). A preservation condition for conditionals. *Analysis*, 60, 219-222.
- Burgess, J. P. (2004). Review of J. Bennett, A philosophical guide to conditionals. *Bulletin of Symbolic Logic*, 10, 565-570.
- Carston, R. (2012). Relevance theory. In G. Russell & D. Graff Fara (Eds.), *The Routledge Companion to Philosophy of Language* (pp. 163-176). New York: Routledge.
- Clark, M. (1971). Ifs and hooks. *Analysis*, 32(2), 33-39.

- Clark, M. (1974). Ifs and hooks: A rejoinder. *Analysis*, 34(3), 77-83.
- Dale, A.J. (1974). A defence of material implication. *Analysis*, 34(3), 91-95.
- Davis, W. (1979). Indicative and subjunctive conditionals. *The Philosophical Review*, 88(4), 544-564.
- DeRose, K. (2002). Assertion, knowledge, and context. *The Philosophical Review*, 111(2), 167-203.
- Douven, I. (2007). On Bradley's preservation condition for conditionals. *Erkenntnis* (1975-), 67(1), 111-118.
- Douven, I. (2008). The evidential support theory of conditionals. *Synthese*, 164(1), 19-44.
- Dudman, V. H. (1988). Indicative and subjunctive. *Analysis*, 48, 113-22.
- Dudman, V.H. (1983). Tense and time in English verb-clusters of the primary pattern. *Australian Journal of Linguistics*, 3, 25-44.
- Dudman, V.H. (1991). Jackson classifying conditionals. *Analysis*, 51(3), 131-136.
- Dunn, J.M. & G. Restall (2002). Relevance logic. In *Handbook of Philosophical Logic, Vol. 6* (pp. 1-128), Springer Netherlands.
- Edgington, D. (1986). Do conditionals have truth conditions? *Crítica: Revista Hispanoamericana de Filosofía*, 18(52), 3-39.
- Edgington, D. (2000). General conditional statements: A reply to Kölbel. *Mind*, 109, 109-116.
- Edgington, D. (2014). Indicative conditionals. *Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/entries/conditionals/>
- Edgington, D. (1995). On conditionals. *Mind*, 104, 235-329.
- Ellis, B. (1969). An epistemological concept of truth. In R. Brown & C.D. Rollins (Eds.), *Contemporary Philosophy in Australia* (pp. 52-72), London: Allen & Unwin.
- Ellis, B. (1978). A unified theory of conditionals. *Journal of Philosophical Logic*, 7, 107-124.
- Ellis, B. (1984). Two theories of indicative conditionals. *Australasian Journal of Philosophy*, 62(1), 50-66.
- Fine, K. (2001). The question of realism. *Philosopher's Imprint*, 1, 1-30.
- Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition and Logical Form*. New York: Academic Press.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121-123.

- Gibbard, A. (1981). Two recent theories of conditionals. In W.L Harper, R. Stalnaker, & C.T. Pearce (Eds.), *Ifs*. Dordrecht: Reidel.
- Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hajek, A. & Hall, N. (1994). The hypothesis of the conditional construal of conditional probability. In E. Eells & B. Skyrms (Eds.), *Probability and Conditionals* (pp. 75-112). Cambridge: Cambridge University Press.
- Jackson, F. (1979). On assertion and indicative conditionals. *Philosophical Review*, 88, 565-589.
- Jackson, F. (1981). Conditionals and possibilities. *Proceedings of the Aristotelian Society*, 81, 125-137.
- Jackson, F. (1987). *Conditionals*. Oxford: Basil Blackwell.
- Jackson, F. (1990). Classifying conditionals. *Analysis*, 50, 134-147.
- Jeffrey, R. (1964). If (abstract). *Journal of Philosophy*, 61, 702-703.
- Kay, P. & Michaelis, L. (2012). Constructional meaning and compositionality. In C. Maienborn, K. von Stechow, & P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning, Vol. 3*, (pp. 2,271-2,296). Berlin: de Gruyter.
- Klein, S. (2013). Lyme disease myths: 9 things you should know about the tick-borne disease. *The Huffington Post*. Retrieved from http://www.huffingtonpost.com/2013/05/20/lyme-disease-myths-facts_n_3287872.html
- Krzyżanowska, K., Wenmackers, S., & Douven, I. (2014). Rethinking Gibbard's riverboat argument. *Studia Logica*, 102, 771-792. Retrieved from <http://karolinakrzyzanowska.com/pdfs/gibbard.pdf>
- Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Levinson, S. (2000) *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, MA: MIT Press.
- Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review*, 85, 297-315.
- Lycan, W. (2001). *Real Conditionals*. Oxford: Oxford University Press.
- Mackie, J.L. (1973). *Truth, Probability and Paradox*. Oxford: Clarendon Press.
- Mayo, B. (1957). Conditional statements. *The Philosophical Review*, 66(3), 291-303.
- McDermott, M. (1996). On the truth conditions of certain 'if'-sentences. *Philosophical Review*, 105, 1-37.

- McGee, V. (1985). A counterexample to modus ponens. *The Journal of Philosophy*, 82, 462-71.
- McGee, V. (1989). Conditional probabilities and compounds of conditionals. *Philosophical Review*, 98, 485-542.
- Mellor, D.H. (1993). How to believe a conditional. *The Journal of Philosophy*, 90(5), 233-248.
- Morgan, J.L. (1978). Two types of convention in indirect speech acts. In P. Cole (Ed.), *Syntax and Semantics, Vol. 9: Pragmatics* (pp. 261-280). New York: Academic Press.
- Nolan, D. (2003). Defending a possible-worlds account of indicative conditionals. *Philosophical Studies*, 116(3), 215-269.
- Numeroff, L. (1985). *If You Give a Mouse a Cookie*. New York: Harper & Row.
- Pendlebury, M. (1989). The projection strategy and the truth conditions of conditional statements. *Mind, New Series*, 98(390), 179-205.
- Quine, W.V.O. (1974). *Methods of Logic*. London: Holt.
- Ramsey, F. P. (1931). General propositions and causality. In D. H. Mellor (Ed.), *Philosophical Papers* (pp. 145-63). Cambridge: Cambridge University Press.
- Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffman (Eds.), *Modality: Metaphysics, Logic, and Epistemology* (pp. 109-135). Oxford: Oxford University Press.
- Russell, L.J. (1960). Formal logic and ordinary language. *Analysis*, 21(2), 25-34.
- Saaed, J. (2009). *Semantics*. Singapore: Wiley-Blackwell.
- Sadock, J.M. (1978). On testing for conversational implicature. In P. Cole (Ed.), *Syntax and Semantics 9: Pragmatics* (pp. 281-298). New York: Academic Press.
- Saenz, N. (2014). The world and truth about what is not. *The Philosophical Quarterly*, 64, 82-98.
- Schaffer, J. (2009). Grounding as the primitive concept of metaphysical structure. Manuscript.
- Schaffer, J. (2014). The Metaphysics of Causation,” *The Stanford Encyclopedia of Philosophy*, <<http://plato.stanford.edu/entries/causation-metaphysics/#Connection>>.
- Sider, T. (2014). *Writing the Book of the World*. Oxford: Oxford University Press.
- Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in Logical Theory, American Philosophical Quarterly Monograph, Series 2* (pp. 98-112). Oxford: Blackwell.
- Stalnaker, R. (1970). Probability and conditionals. *Philosophy of Science*, 37, 64-80.

- Stalnaker, R. (1975). Indicative conditionals. *Philosophia*, 5, 269-86.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, R. (2005). Conditional propositions and conditional assertions. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality* (pp. 227-248). Oxford: Oxford University Press
- Stanley, J. (2002). Modality and what is said. *Philosophical Perspectives*, 16, 321-344.
- Stevenson, C. (1970). If-iculties. *Philosophy of Science*, 37(1), 27-49.
- Strawson, P. F. (1981). 'If' and '⊃'. In R. Grandy & R. Warner (Eds.), *Philosophical Grounds of Rationality: Intentions, Categories, Ends* (pp. 229-242). Oxford: Oxford University Press.
- Sweetser, E. (1991). *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge: Cambridge University Press.
- Tichy, P. (1976). A counter-example to the Stalnaker-Lewis account of counterfactuals. *Philosophical Studies*, 29, 271-273.
- Tolkien, J.R.R. (1994). *The Hobbit*. Boston, MA: Houghton Mifflin.
- Weatherson, B. (2001). Indicatives and subjunctives. *Philosophical Quarterly*, 51, 200-216.
- Weatherson, B. (2014, April 15). Conditionals, predicates, and probability. Retrieved from <http://brian.weatherson.org/adams.pdf>
- Young, J. (1972). Ifs and hooks: A defence of the orthodox view. *Analysis*, 33(2), 56-63.

APPENDIX

I. APPENDIX ONE: Proof that the probability of B, given A, is equal to B in cases in which the probabilities of A and B are independent.

$$\Pr(B | A) = \Pr(A \& B) / \Pr(A)$$

$$\Pr(A \& B), \text{ where } \Pr(A) \text{ is independent of } \Pr(B), = \Pr(A) \cdot \Pr(B)$$

$$\text{So, if } \Pr(B) \text{ is independent of } \Pr(A), \text{ then } \Pr(B | A) = \Pr(A) \cdot \Pr(B) / \Pr(A) = \Pr(B)$$

So, if $\Pr(B)$ is high, as is the case in every BB conditional, then $\Pr(B | A)$ is high.

II. APPENDIX TWO: Case in which $\Pr(A)$ and $\Pr(B)$ are not independent, $\Pr(A)$ is high, $\Pr(B)$ is high, and $\Pr(B | A)$ is low.

Suppose you have three dice: yellow, red, and green. You roll the yellow die first. If the die comes up 1-4, you then roll the red die (so there is a $2/3$ chance that you roll the red die). If the yellow die comes up 5-6, you roll the green die. Let A be 'you roll the red die' and B be 'you roll a red 1 or 2, or you roll the green die'. In this case,

$$\Pr(A) = 2/3, \text{ which is reasonably high.}$$

$$\Pr(B) = \Pr(B \& A) + \Pr(B \& \sim A) = (2/3 \cdot 1/3) + (1/3 \cdot 1) = (2/9 + 3/9) = 5/9, \text{ which is somewhat high.}$$

$$\Pr(B | A) = 1/3, \text{ which is low.}$$