Reasonable Assertions:
On Norms of Assertion and Why You Don't Need to
Know What You're Talking About

by

Rachel McKinnon

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Philosophy

Waterloo, Ontario, Canada, 2012

AUTHOR'S DECLARATION FOR ELECTRONIC SUBMISSION OF A THESIS

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Rachel McKinnon

Abstract

Theres a widespread conviction in the norms of assertion literature that an agents asserting something false merits criticism. As Williamson puts it, asserting something false is likened to cheating at the game of assertion. Most writers on the topic have consequently proposed factive norms of assertion ones on which truth is a necessary condition for the proper performance of an assertion. However, I argue that this view is mistaken. I suggest that we can illuminate the error by introducing a theoretical distinction between the norm of a practice and its goal. In light of this distinction, we can see that proponents of factive norms tend to mistake the goal of a practice for the norm. In making my case, I present an analogy between the norms and goals of placing wagers and the norms and goals of assertion. One may place a bet and lose without being subject to criticism, while one may win and be worthy of criticism. Whether one wins or loses is irrelevant to the normative evaluation of a bet. What is relevant is whether the bet maximizes the bettor's expected value, which is a function of what might be lost, what might be gained, and how likely those prospects are, given the bettor's evidence. Similarly, I argue, whether one's assertion is true or false is not strictly relevant to the normative evaluation of an assertion. What is relevant is whether the speaker has adequate supporting reasons for the assertion, and that the necessary conventional and pragmatic features are present. However, context will determine what count as supportive reasons for a given proposition, what counts as relevant, and what count as conventional and pragmatic elements possessing that relevance. My proposed norm, the *Supportive Reasons Norm*, is thus sensitive to the context of assertion and shifts from context to context.

Acknowledgements

First and foremost I'd like to thank Tim Kenyon for his guidance, patience, and support over these past few years. His mentorship has made me a better philosopher, writer, teacher, and colleague. Dave DeVidi has also been a formative force in my career, always willing to add just a little bit more work to the giant pile on his desk, or find some time for some helpful advice. He's also the one who most keeps me on my toes, sometimes with only a single raised eyebrow. I'd also like to thank John Turri for his many conversations and comments on my work. I've tried to approximate his exceptional work ethic and ability to succinctly represent an argument, while injecting a little well-chosen humour.

I'd like to thank my many friends and colleagues. Their support and friendship made this possible, and some were even unfortunate enough to offer to provide comments on drafts of some of my work. My parents have also played an important supporting role, especially by providing opportunities to poke my head up from my work now and then. I know he can't read this, but I'd also like to thank my dog, Dennett. It's simply impossible for me to stay in a foul mood with him around. He kept me sane and happy. I love you buddy, more than you can probably know.

Table of Contents

Chapter 1

Introduction

1.1 Assertion Within Reason

A recent area of interest for epistemologists and philosophers of language has been the question of what norms, if any, govern or characterize the linguistic practice of assertion. Assertion is a central speech act in our communication, and it's so ubiquitous as to often go without notice or critical reflection. We find assertions everywhere: in giving directions, describing one's future plans, responding to questions, expressing one's political views, and in many other places. But why think that assertion is a norm-governed practice? One reason is that there are many intuitive senses in which we wish to criticize various assertions. A lie is a clear case of a faulty assertion: we blame the speaker for having done something wrong.

However, there may be other less obviously faulty assertions. Suppose that I see a memo for a departmental meeting. The memo indicates that the meeting is at 3:30 p.m., which is the usual time. A new colleague, Ben, asks me what time the meeting begins, to which

I respond, "3:30 p.m." But suppose that, unbeknownst to me, the meeting time was just changed to 4:00 p.m., and notice has yet to go out about the change. So my assertion is false, but did I do anything wrong in performing that assertion? In a sense, if I were to find out about the time change, there's an intuition that I ought to retract my assertion that the meeting is at 3:30 p.m. However, on the other hand, there's an intuition that although what I said is false, I can't be blamed for it: I had excellent evidence for the meeting being at 3:30 p.m., and I had no reason to think that it's at 4:00 p.m. But the question remains: was my assertion, all things considered, good, or not? This question strikes at the heart of the norms of assertion debate. Advocates of the knowledge norm (one should assert that p only if one knows that p) argue that the assertion is inappropriate, but I'm not deserving of blame. They make a distinction between criticizing and blaming. Advocates of what I call reasons-based norms (*e.g.*, assert that p only if it's reasonable for one to believe that p) argue that the assertion is appropriate, even praiseworthy, though it would have been better if the assertion were true.

In considering these issues, we raise a number of important questions. For example, what would such norms of assertion look like? What is the connection between assertion and epistemic concepts such as truth, justification, and knowledge? Are there some propositions that one could know but not warrantedly assert? What sorts of things do we do with assertions? What is the assertoric status of bullshit? The purpose of this dissertation is to take up these issues in the service of articulating norms of assertion.

One aim of this dissertation is to sort through a number of the potentially conflicting

intuitions surrounding assertion, with the aim of providing a unifying analysis. There's a wide range of disparate intuitions marshalled for and against various views on the norms of assertion. So it's important for any new proposal to attempt to offer a unifying account. Such a project will require a number of detours following various interrelated threads of inquiry. For example, this will include discussions about what it means to "aim" at truth and whether a deflationary theory of truth is adequate to fill truth's role in assertoric norms; whether statements about lotteries, made prior to information of the results of the draw, are knowable or warrantedly assertible; what evidentiary role challenges and prompts to assertions, and how one ought to respond, have for assertoric norms; and how assertion can be modelled on how professional gamblers ought to reason in placing bets.

1.2   Brief Overview of the Literature

Most candidate norms focus on epistemic concepts like truth, reasonable belief or justification, and knowledge. There are three prominent proposals for the central norm of assertion. These are the Truth Norm (TNA), the Knowledge Norm (KNA), and some version of a Reasonable Belief Norm of Assertion (RBNA). Other less prominent proposals include the Belief Norm (BNA) and the Certainty Norm (CNA). Roughly speaking, versions of TNA hold that an agent may assert only if the proposition asserted is true; versions of KNA hold that an agent may assert only if she knows the proposition asserted; and versions of RBNA hold that an agent may assert only if she reasonably believes the proposition asserted.

Although there is a great deal of focus on the epistemic requirements for warranted

assertibility, there are other factors relevant to the nature and content of assertoric norms. Assertion is a public practice: it's something done with other competent speakers. It doesn't make much sense to assert to a tree, after all. Consequently, there are a number of social conventions surrounding the appropriateness of some assertions. Although an assertion may be true, and even a case of knowledge, it may be impolite to assert during a department meeting that your boss has very bad breath. So there may be important conventional aspects to assertoric norms. Additionally, there may be important pragmatic aspects relevant to assertibility. There are many things that we do with assertions that may bear on assertibility conditions. For example, assertions are typically taken as reasons for action. Although there may be a picture of fire on the wall of a crowded movie theatre, it would be a rather terrible idea to suddenly loudly assert "Fire!" People will wrongly interpret this true assertion as a reason to stampede for the exits (naturally thinking that there's a real fire in the theatre). Since this predictably bad outcome is a likely consequence of the assertion, it makes the assertion inappropriate. Understanding the pragmatic implications of assertion, including what conditions are required for licensing an interlocutor to take one's assertion and act on it, might therefore bear on whether one may or ought to assert.

The literature on assertoric norms has focused on two principal broad categories of data and arguments. The first involves conceptual arguments concerning the nature of assertion as a speech act distinct from, say, questioning, conjecturing, or commanding. The more recent literature, however, has heavily focused on a second class of data which we may call the "linguistic data." This involves a wide class of observations and intuitions about the

4

everyday uses of assertion. We may separate the linguistic data into three broad categories: Moore's Paradox, which involves statements such as "Today is Friday, but I don't know that today is Friday"; the Lottery Paradox, which involves statements about lottery tickets made prior to information on the results of the draw such as "Your ticket will lose"; and additional observations about topics such as the propriety of various challenges to assertions such as "How do you know?" This latter data involves articulating both how a speaker ought to respond to both prompts (such as "What time does the meeting start?") and challenges, in addition to how hearers respond to various assertions.

Throughout the literature, arguments for proposed norms of assertion typically take the form of *inferences to the best explanation.* Some argue that a particular norm best explains some subset of the data, and that this provides evidence for such a norm. For example, Timothy Williamson has argued that the KNA is best able to explain the Lottery Paradox and that this supports KNA.[1] Others argue that a particular norm best explains all of the data, when taken as a whole, compared to other candidate norms. For example, John Turri has argued that KNA can best account for all of the linguistic data of challenges and prompts, even though some competitors may do better on subsets of the data, and that this supports KNA.[2] I will have more to say about a good methodology for pursuing norms of assertion in Section 1.4 below.

---

[1]Williamson (2000).
[2]Turri (2010c) and Turri (2011).

1.3 The Supportive Reasons Norm of Assertion (SRNA)

In this dissertation I will defend a view of assertoric norms on which one may assert only if one has adequate epistemic reasons for the assertion that p for the context, and the necessary conventional and pragmatic considerations are present.[3] For simplicity's sake, I will call this the *Supportive Reasons Norm of Assertion* (SRNA). In an important sense this is a reasons-based norm; yet it differs significantly from extant reasons-based norms, including Jennifer Lackey's Reasonable to Believe Norm of Assertion (RTBNA).[4] One of the biggest differences is how I incorporate the conventional and pragmatic features of assertion *into* the norm rather than attach them as an external addendum. SRNA also takes seriously the observation that assertibility seems to change when the practical stakes of the context of assertion changes. For example, when the stakes are higher, better reasons are required. However, there are even some contexts where a speaker may warrantedly assert something she knows to be false, such as some pedagogical contexts. As such, there currently are no other proposals in the literature for an assertoric norm quite like SRNA.

Calling it SRNA and roughly defining it as I have belies its complexity. I will argue that the *central* epistemic concept to assertion is epistemic justification, evidence, or what I'll simply call *reasons*. (I'll use "justification," "evidence," and "reasons" interchangeably, but I'll generally use "reasons.") By this I mean that there are no cases of warranted assertions where adequate reasons, in an important epistemic sense, are lacking. However, depending

---

[3]Though what I take to be "adequate epistemic reasons" is highly specific and, possibly, non-standard. I will explain below.

[4]Douven (2006), Hill & Schechter (2007), and Kvanvig (2009) are other reasons-based norm advocates.

on the conventional or pragmatic considerations of a particular case, the speaker herself need not supply the justification. In such cases it may be sufficient for the asserted proposition to be a justifiably believable proposition for the hearer alone. But justification must be present either on the part of the speaker or the hearer.[5] For example, if a speaker knows that something is false, then she cannot have epistemic justification for the proposition. However, she may warrantedly assert the false proposition if the proposition is something that could be justifiably believed by the hearer. Cases such as this are not hard to imagine: consider the wide range of pedagogical assertions made by teachers which are, strictly speaking, known to be false but it's best for the student to be told a noble lie rather than the truth. It's better to instruct young students in physics by using the Bohr model of the electron structure of atoms rather than the valence model, for example.

Another important feature of SRNA is that the speaker need not believe what she asserts. Cases of warranted assertions will often involve the speaker believing what is asserted, but, provided that the appropriate conventional and pragmatic considerations are present, this is not a necessary condition. The pedagogical example above is such a case, as are Lackey's cases of "selfless assertions."[6] A further defining feature of my account is the importance given to the highly variegated conventional and pragmatic features of assertion. This raises natural questions such as, "What determines which conventional and pragmatic features are relevant?" and, "How do the 'relevant' conventional and pragmatic features impact the

---

[5]The speaker knowing that a proposition is false would not count as a defeater to the hearer's justification. The speaker can be a source of epistemic justification for the hearer on broadly testimonial grounds even though, in such a case, the hearer cannot come to *know* what is asserted because the proposition is false. I also leave open the possibility that the supportive reasons for an assertion lay with a third party (*i.e.*, with neither the speaker or the immediate hearer).

[6]Lackey (2007).

requirements for warranted assertibility in a given case?" The short answer is, "It depends."

Assertion is not a simple practice. Much like other complicated social practices, such as

etiquette, law, and government, a one-size-fits-all approach will not adequately capture the

phenomena.[7] It is a primary task of my dissertation to sort through many of these issues.

This context-sensitivity of my account connects with the relevant literature on epistemic

contextualism in the following respects. First, my position on warranted assertibility is

invariantist in that the requirements for warranted assertibility do not shift with context:

agents must always have adequate evidence (however conceived for the case) and the relevant

(contextual) conventional and pragmatic features must be present in order to warrantedly

assert some proposition. However, the satisfaction of those conditions is highly sensitive to

context. So the truth conditions for ascribing the predicate "warrantedly assertible" *will* shift

with context, but my definition of what conditions are required for warranted assertibility

will not.

Although I argue that justification is the central epistemic concept in assertion, the con-

textualism of my theory allows it to deal with cases where circumstances clearly seem to

require speakers to have particularly strong epistemic support for what they assert, such as

in high-stakes court testimony. Although my theory argues that the asserted proposition

being a suitable candidate for reasonable belief is the guiding epistemic requirement, the

conventional and pragmatic features of testifying in high-stakes court cases means that one

must have particularly strong reasons in order to properly assert. KNA advocates, by con-

trast, have a difficult time explaining why some circumstances licence assertions when the

---

[7]Insofar as both assertion and law are *practices*, the analogy is appropriate.

speaker lacks knowledge, but has good reasons. That is, some circumstances may require an expedient assertion even when one has doubts as to the truth of the proposition asserted. For example, I may assert, "That's your train. Run!" if missing the train is dire, but being wrong is not. I will argue that my account best explains all of the forms of assertoric data when taken as a whole, and explains each of the data at least as well as any competing theory when taken individually.

1.4   Methodology

There are a number of conceptual issues to be settled regarding articulating norms of assertion. What types of norms might assertoric norms be? What is the connection between truth and assertion? How should we properly interpret the "data?" There are also empirical issues to be settled. How do we actually engage in the practice of assertion? How do people react to assertions that are false? Are there incorrect assertions? So the project of articulating possible assertoric norms is at the same time conceptual and empirical, *a priori* and *a posteriori*. But there are also methodological issues to be settled in how one ought to approach this project. Any theory of assertoric norms should be sensitive to "save the phenomena." That is, theories should at least explain the important intuitions and linguistic observations. There is a *prima facie* requirement to explain the various linguistic data as they are. But this is not to say that all observations are created equal and should take precedence over strong theoretical reasons for re-interpreting the data. For example, in Chapter 7 I will offer a re-interpretation of the linguistic data of challenges such as "How do you know?"

This challenge is typically considered as challenging a speaker's knowledge, but I will suggest that there are strong theoretical and empirical grounds to doubt this. It should be possible, therefore, that our intuitions can be mistaken and should be reconsidered in light of strong theoretical considerations.

A common methodology in the literature is to seek a single, simple constitutive norm of assertion. Although a more detailed discussion will be saved for Chapter 3, a norm is constitutive of a practice if, necessarily, every act of that type must satisfy the norm in order to be appropriate. Furthermore, constitutive norms "create" the practice and may be said to be "essential" to the act. For Williamson, "if it is a constitutive rule that one must ϕ, then it is necessary that one must ϕ."[8] Constitutive norms are thus necessary, but not sufficient, conditions for properly performing an act in a practice. Williamson's proposed norm of assertion is the Knowledge Norm (KNA) on which one may (must) assert some proposition, p, only if one knows p.[9] This is not yet to provide sufficient conditions, but it does mean that any assertion of a proposition that a speaker does not know is an improper assertion.[10] KNA is also a "simple" norm in that a single rule of *assertion* is proposed. There may be other norms that contribute to the various sufficient conditions required for proper assertion, but they either will not be unique or constitutive of assertion, or will be derivative from KNA. Such norms may be, for example, general conversational (*e.g.*, Gricean) or moral norms such as "Be truthful," or "Be polite."

---

[8]Williamson (2000, p. 239).

[9]Some may go so far as to offer a norm of the form: Assert that p if and only if one knows that p.

[10]There are some exceptions for Williamson, particularly in cases where we may be "lax'" concerning the enforcement of the knowledge norm. However, it remains that knowledge *is* the norm of assertion.

There remains a question concerning how we should treat theoretical considerations such as the simplicity of a theory when compared to another. Often in science, when two competing theories or hypotheses can equally explain the data, the simpler theory is favoured. Some, such as Igor Douven, have suggested a similar tactic for assertoric norms: if KNA and some competitor norm are equally able to explain the linguistic data, then we should favour the simpler norm.[11] However, what reasons do we have to think that assertion is a simple practice describable by a single, simple constitutive norm? A number of philosophers have likened the norms of assertion to rules of a game.[12] If the analogy is close, then we would expect many of our games to be governed by simple rules. But this is not what we find. The games of baseball, basketball, hockey, football, soccer, tennis, badminton, golf, pool, darts, and so on, all have long and complicated rule-books. Even chess and checkers are not governed by single, simple constitutive rules.

Perhaps our social practices are not strongly analogous with games, but perhaps a practice like assertion may still be governed by a single, simple constitutive norm. But this is not what we seem to find in other social practices and institutions. Look to our social institutions of government and law, or social practices such as driving, being a good colleague, and having good manners. Insofar as any of these can be a distinct practice, it's rather clear that there are many norms, often in conflict, and which can be quite complicated. So rather than having good reasons to think that assertion would be governed by a single, simple constitutive norm

---

[11]Douven (2006). He uses this strategy to argue for his "rational credibility account of assertion'" and the *Rational Credibility Rule*.

[12]Dummett (1981), Williamson (2000), Lackey (2007), and others. See Maitra (2011) for an opposing view.

– assertion is "a part of reality of our making" after all[13] – we have strong reasons to think that assertion would *not* be governed by such a norm. Thus, I reject the Williamsonian (*et al.*) strategy of assuming that assertion is governed by a single, simple constitutive norm. It's not clear what methodological advantages such a strategy provides.

Another question of methodology concerns how the argumentation will proceed. As already indicated, many arguments in the literature take the form of an *inference to the best explanation* (IBE). For example, upon canvassing all of the conceptual and linguistic data of assertion, one might argue that KNA can *best* explain the data when taken as a whole. Perhaps KNA has some trouble explaining cases where assertions are appropriate in the absence of knowledge, but this may be fixed by appeal to how the knowledge requirement might be overridden by extenuating circumstances in special cases. In such cases a reasons-based norm might appear to better explain the data, but this is an isolated segment of the data. So when the data is taken *as a whole* KNA might do better than any of its competitors. This is one form of an IBE. Another would also require that a candidate norm explain each subset of the data *at least as well* as any competitor, while explaining the data, taken as a whole, better than any competitor. All things being equal, if a norm can best explain the data when the data is taken as a whole, but can also explain each of the subsets of data when the data is taken individually, at least as well as any competitor, then such a norm appears to be well supported both theoretically and empirically. This is the argumentation strategy that I will adopt in my argument in support of SRNA.

There is also a methodological tension between simplicity and completeness. I will adopt

---

[13]Douven (2006, p. 451).

the methodology of favouring an account that is, *ceteris paribus*, the most simple. However, as I will argue, the *ceteris paribus* requirement is rarely satisfied, and when all of the relevant considerations are included, SRNA will be the most complete. I will argue that competing accounts such as TNA, KNA, and RTBNA come out of the analysis with considerable bruises or having suffered damage due to apparently decisive counterexamples. In accomplishing this goal, SRNA incorporates the complexity of the practice of assertion, but not on an *ad hoc* basis. In the end, I argue that many of the KNA treatments of the data depend on a reasons-based norm explanation. Consequently, not only does KNA seem to get some important cases wrong, positing KNA is superfluous for accounting for the data.

There's a further methodological question regarding how one ought to construct a proposed norm of assertion. The question is: do we take the linguistic data and then conjecture an explanatory norm that fits the data, or do we start from considerations about the nature of assertion and conjecture a norm which is then tested by the data? It appears that most of the literature on norms of assertion has adopted the former approach.[14] We observe that some cases of assertion without knowing what one asserts appear inappropriate; therefore we posit KNA. But we notice that some cases without knowledge *are* appropriate, but these are cases where truth is still required; therefore we posit TNA. Then we notice that there are still some cases where neither knowledge nor truth are required, and that in every case one must assert only if it's reasonable for one to believe what one asserts; therefore we posit RTBNA. But I think that this programme has things backwards: we should start from the-

---

[14]Of course, I say "most" instead of "all" because some argue that the aim of assertion is to express knowledge and that, therefore, knowledge is the norm of assertion. In Chapters 4 and 8 I argue that this is a poor argument. I argue that this conflates the goal of assertion with the norm.

oretical considerations of the nature of assertion, understanding its constitutive aims, how norms relate to these aims, and then posit a norm which seems best suited to fill that role. Then we take the various data and test the proposed norm.

This is how I will argue for a reasons-based norm: not by taking the data as a whole and positing a norm which best explains the data; rather, I will begin by considering the constitutive aim(s) of assertion, what it means for something to be a norm, and then conjecture a version of a reasons-based norm. In Chapter 8 I propose just such a norm, SRNA, and then in Chapter 9 I test the norm against the data. I will show that SRNA performs at least as well as any other norm on any given subset of the data, such as KNA in explaining the lottery paradox data, and that SRNA explains the data as a whole better than any other norm. This puts SRNA on better theoretical grounds, I think, in addition to the empirical grounding provided by explaining the data.

1.5   Summary of Chapters

Building a case for SRNA will require some preliminary discussions pertaining to understanding the nature of both assertion and norms. In Chapter 2 I begin with a discussion of the nature of the speech act of assertion. There are a number of options in characterizing assertion as a speech act distinct from, say, commanding or conjecturing. One definition simply states that assertion is a speech act where some proposition is claimed to be true, *e.g.* that today is Friday. I think that this is correct, as far as it goes. Additionally, though, I am partial to defining assertion in terms of the conventional commitments that a speaker

acquires in asserting some proposition. Moreover, like Brandom, I favour a definition on which assertion constitutes a speaker's giving a hearer *reasons* for believing the proposition asserted.

In Chapter 3 I will discuss topics related to specifying the concept and content of norms. This will include specifying the differences between constitutive and non-constitutive (*e.g.*, regulative), simple and complex, restrictive and non-restrictive, and defeasible and indefeasible norms. It will also discuss how some norms are clearly action-guiding, whereas others merely specify the relevant goals of action (without thereby providing guidance in how one ought to act to obtain that end). I will also discuss an often overlooked topic: the relationship between a *goal* and a *norm*. I'll argue that norms derive their content, in part, from what it takes to form rational plans for achieving the goals of a practice. I'll suggest that mistaking the relationship, and necessary separateness, between norms and goals of a practice is a key source of errors in the literature on norms of assertion. This is a distinction that I'll put to work throughout the dissertation in building my case for a reasons-based norm.

Chapter 4 will begin by addressing the intuitive connection between proper assertion and truth. So I begin by taking on what I consider to be truth's role in assertoric norms. I will argue that warrantedly asserting is "properly aiming" at truth, but actually obtaining truth is not a necessary condition. This will involve discussing accounts which tie the meaning of truth to warranted assertibility in terms of *endorsing* assertions. We endorse assertions, for example, when we say of an assertion, "That's true." This has implications for theories of truth; namely, whether a deflationary theory of truth is adequate for truth's role in assertion.

In Chapter 4 I also directly discuss the arguments in favour of one of the primary candidate norms of assertion: the Truth Norm (TNA). There is a strong intuitive connection between proper assertion and truth. This makes TNA a strong *prima facie* candidate for the central norm of assertion. However, I will argue that there are troubling cases for TNA. I will focus on two: lucky guesses and the lottery paradox, though the latter will be more fully developed in Chapter 5.

In Chapter 5 I will present the most popular assertoric norm in the literature, and the biggest competitor to my account: the Knowledge Norm (KNA). Chapters 5 through 7 constitute a campaign against the most common arguments in favour of KNA: the lottery paradox, Moore's paradox, and challenges to assertion, respectively. Although, Chapter 8 will feature another line of argument offered in favour of KNA. A common argument supporting KNA is the intuitive impropriety of making assertions about lottery tickets prior to information on the results of the draw. Even though a speaker may assert something that has an arbitrarily high probability of being true on the available evidence (and the assertion may actually be true), such assertions seem to be categorically inappropriate. KNA proponents, such as Williamson, argue that this is explained by the speaker not *knowing* what she asserts. However, I will argue such accounts are unsatisfactory. I will argue that a relevant alternatives approach will best solve the problem in the literature on why one can know normal propositions (such as whether I am presently typing at a computer) but one cannot know propositions about the outcomes of lotteries made prior to the results of the draw (such as whether my ticket will lose). I argue that the pathology with such

assertions, and beliefs, lies in that one does not have adequate epistemic evidence for such propositions. This explains why speakers do not know what they assert: because they lack adequate evidence. But identifying the pathology of such assertions this way removes the advantage the KNA advocate claimed to have: now other norms, particularly reasons-based norms, are easily able to account for the data. So, I argue, the lottery paradox data doesn't uniquely support KNA rather than a reasons-based norm or, possibly, even TNA.

Chapter 6 continues to take on what has broadly become known as the "linguistic data" of assertion. Other than the assertibility of lottery propositions ("Your ticket will lose"), another general class of data is known as Moore's (Knowledge) Paradox and concerns the intuitive oddity of statements such as "Today is Friday, but I don't know that today is Friday." There has been considerable focus on how various theories of assertoric norms can explain the apparent impropriety of such assertions. However, I will argue that this focus is misplaced. I will argue that all of the proposed assertoric norms are able to explain Moore's Paradox and, consequently, Moore's Paradox does not provide any differential support to one candidate norm over another. That is, Moore's Paradox does not stand as evidence for any given norm; instead, it should be treated as an adequacy condition of assertoric norms.

Chapter 7 attempts to remove yet another pillar of support in the case for KNA by focusing on the general class of linguistic data involving the propriety of various prompts – such as "What time does the meeting start?" – and challenges – such as "How do you know?" – to assertions and how speakers ought to respond to such prompts or challenges. Some, such as Turri, have argued that the default propriety of challenges to a speaker's *knowledge*

provides particular support for KNA.[15] However, I will argue that the data is not currently properly interpreted. I will offer a unifying analysis of this data such that we can understand the content of a challenge or prompt based on what constitutes a *wholly adequate response* on the part of the speaker. I will argue that the propriety of various forms of prompts and challenges, and their attendant responses, is best explained by a reasons-based norm (such as RTBNA or SRNA) rather than KNA.

In Chapter 8 I will present my case for SRNA. I will begin by offering a new model of assertoric norms based on the norms of properly placing wagers and responding to the results. The critical result will be that, in many circumstances, the results of a wager do not bear on its propriety. That is, whether a wager wins or loses does not bear on whether an agent made a good bet and, moreover, on whether the agent is fully praise- or blameworthy. I argue that this supports a non-factive norm (which eliminates KNA, TNA, and CNA), and seems to support a reasons-based norm. I will then use a recent new attempt to support KNA by Ernest Sosa in explicit terms of performance normativity and aiming at goals, both concepts I consider central to understanding norms of assertion.[16] Since he uses these same central concepts in arriving at KNA rather than a reasons-based norm, I will use his view as a foil for my own. I will argue that Sosa's account makes the same mistake as the TNA advocate in equating the *norm* with the (plausible) *goal* of the practice.

Finally, in Chapter 9 I conclude the dissertation by taking stock. The principal goal of this chapter is to test SRNA, synthesizing the arguments for reasons-based norms in the

---

[15]Turri (2010a).
[16]Sosa (2011).

preceding chapters. I will explain how it is able to properly account for each subclass of the data at least as well as any competitor norm and is able to *best* explain all of the data when taken as a whole. I then summarize the various candidate norms and how they stack up against my own account. I end with some discussion of directions for future research. For example, by more fully understanding the connection between warranted assertion and central epistemic concepts such as knowledge, truth, justification (and so on), we may come to better understand these epistemic concepts. That is, what we *do* with epistemic concepts may tell us something important about their nature. So I will point to promising directions of research into what we *do* with knowledge such as practical decision-making and discussions on pragmatic encroachment in how we understand knowledge.

Chapter 2

Assertion as a Speech Act

2.1 Introduction

Asserting is one of many acts we can perform in uttering words. We find assertions everywhere in our daily discourse when we give directions, answer questions, offer our opinions, and when we make arguments. But what is an assertion and what separates it from other speech acts such as requesting, conjecturing, and commanding?[1] I will argue for a characterization of assertion in terms of the commitments the speaker takes on by putting forward some proposition as true. Moreover, we should understand both what it means to put forward a proposition as true and speaker commitments in terms of the social and conventional aspects of the practice of assertion. This is a position that has antecedents in the work of Dummett and Brandom, though my own presentation may be subtly different in some respects. One of

---

[1] For a very good, recent survey of assertion and the various theories (including the author's own), see Jary (2010).

the elements or corollaries of this view of assertion is that we shouldn't expect sharply defined boundaries to the phenomenon. But understanding assertion in this way will at least give us purchase, in Chapter 3, on the project of understanding assertion through understanding its norms.

I begin by considering some features of assertion found in speech act theory. I then consider Dummett's influential view of assertion. I complement Dummett's view with Brandom's, making explicit which aspects of their views I will retain and which I will reject. I will then be in a position to state my view of assertion, which is a composite of Dummett and Brandom's views.[2] I finish by introducing an additional methodology for understanding assertion by attempting to articulate its norms, which will be taken up in more detail in Chapter 3.

2.2   Linguistic Acts and Speech Acts

Assertion is one of the many forms of speech acts known as an *illocutionary* act.[3] Illocutionary acts are what we *do* with utterances such as when we assert, command, request, and so on. For example, "Can you close the door?", depending on the context, may either be a request to close the door, a command to close the door, or a question concerning whether the hearer is capable of closing the door. We can contrast illocutionary acts with *locutionary* acts. (Alternatively, one may call locutionary acts *linguistic* acts, and illocutionary acts *speech*

---

[2]For the sake of brevity, I will not be discussing, at any considerable length, rival accounts such as Searle (1969) and Searle (1979) or Bach & Harnish (1979).

[3]The distinction comes from Austin (1962).

acts.) A locutionary act is simply the performance of an utterance, such as "Can you close the door?" Thus, the same locution can comprise many different illocutions, depending on the context.

Consider the following cases:[4]

1. A: Have you been in Boston before?

   B: Yes.

2. A: Do you promise that you'll come to my party?

   B: Yes.

3. A: Do you apologize for your rude behaviour?

   B: Yes.

In each case an act of the same utterance type, saying "Yes," effects a different speech act, depending on the context given by the different questions asked of B; asserting in 1, promising in 2, and apologizing in 3. Following John Turri, "let *speech act contextualism* be the view that context affects which speech act you perform in performing a linguistic act."[5]

While it's generally uncontroversial that which speech act one performs with a given linguistic act depends on context, some have argued for speech act *pluralism*: that, *e.g.*, in asserting "Today is Tuesday" I thereby also assert that it is not Wednesday and a variety of other propositions entailed by the proposition asserted. Cappelen & Lepore (2005), for example, argue that "[n]o one thing is said (or asserted, or claimed, or . . .) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, stated. What is said

---

[4]Taken from Turri (2010c). Compare Fogelin & Sinnott-Armstrong (1997, p. 10).
[5]Turri (2010c, p. 10). Emphasis in the original, but I have changed bold to italics.

(asserted, claimed, *etc.*) depends on a wide range of facts other than the proposition semantically expressed. It depends on a potentially indefinite number of features of the context of utterance and the context of those who report (or think about) what was said by the utterance."[6] So in performing a speech act one may be performing multiple simultaneous speech acts.

On its face, this is contentious: while many propositions may be *deduced* from what is asserted, it's harder to find it plausible that a speaker really expresses all of these propositions by the single utterance. Although by asserting "Today is Tuesday" one can thereby deduce that I must also mean, or be rationally committed to, the further claims that, "Today is not Wednesday," "Today is not Thursday," "Today is not Friday," and so on, it's quite another thing to suggest that I express all of these indefinitely many further propositions by the single assertion, "Today is Tuesday."

In another sense, however, speech act pluralism is not altogether controversial: we regularly recognize many situations where a single utterance performs multiple speech acts. Suppose that "swearing" as found in legal testimony – "Do you hereby swear that your testimony was the truth?" – is a distinct speech act from asserting. In answering "Yes, my testimony was the truth," a speaker may be swearing *by* asserting that her testimony was the truth. So we could understand speech act pluralism in terms of a series of by-relations in that one performs one speech act *by* performing another speech act.

However, we could also understand speech act pluralism as the position that a single utterance can perform multiple speech acts without clear by- relations. Suppose that James is

---
[6]Cappelen & Lepore (2005, p. 4).

a clever FBI double-agent working deep undercover in a mafia organization. He is constantly under close scrutiny by his mafia compatriots, and he is also under constant surveillance by his government allies. There are a number of pre-set phrases that James can use to communicate to the FBI, such as, "That hit was the most masterful performance since Nolan Ryan's last no-hitter." (By "hit" I mean "execution.") Suppose that this is interpreted by the FBI agents as saying, "The mafia boss is in my presence and you should move in to capture him." In a number of contexts, it will be appropriate for James to praise a mafia execution with this linguistic act, and he'll be asserting that the execution was well performed. It's important, after all, that when James is trying to communicate with the FBI that his utterances make sense to his mafia compatriots, so it would be out of place for him to utter this if it wasn't plausible that he be making an assertion here. But since this phrase is also code for the FBI, James will be simultaneously making a completely different assertion to the FBI: the quality of the execution is irrelevant and all the FBI cares about is that the mafia boss is in James's presence.

One might object that in some contexts it would be best to describe James's uttering "That hit was the most masterful performance since Nolan Ryan's last no-hitter" as *not* making an assertion about the quality of the hit, but only as the assertion to the FBI that the mafia boss is in James' presence. However, we can imagine at least some contexts where the utterance really is doing double duty: as an assertion both about the quality of the hit *and* about the presence of the mafia boss. In such cases James will be communicating different propositions to the mafia and the FBI.

24

2.3 Dummett on Assertion

A natural way of distinguishing assertion from other speech acts, such as commanding and requesting, is understanding assertion as the external expression of our internal judgment that some claim or proposition is true. This approach individuates the linguistic act by uniquely linking it to a mental act that is itself already (putatively) clearly individuated. If I say, "It's cold today," this is naturally understood as my giving voice to my belief that it's cold today.[7] As Timothy Williamson has claimed, "believing p stands to asserting p as the inner stands to the outer."[8] Others have gone further to say that assertion is, at its core, the *expression* of a belief.[9] This is partially borne out in our intuitive understanding of assertion as a primary means of communicating information: how do I tell you what I think about some issue? I assert!

John Searle has an account of speech acts on which each illocutionary act (*e.g.*, promising, asserting, requesting) has a specific *structure*.[10] All illocutionary acts have *preparatory*, *sincerity*, and *essential* conditions. Preparatory conditions are necessary conditions for an act to be non-defective in a context and in order for the act to have a point. Thus, in a sense, failing to meet preparatory conditions produces a self-defeating act. Giving an order, for example, has the preparatory condition that the speaker is in a position of authority over the hearer: a non-defective order can't take place unless this is the case. Sincerity conditions

---

[7]I have transitioned from talk of judgement to the mental state of belief. This is due, in part, to Brandom's (and Dummett's) reading of Frege on how an assertion is the giving voice to a thought. I will use "judgement" and "belief" interchangeably unless otherwise noted.

[8]Williamson (2000, p. 255).

[9]See, for example, Bach & Harnish (1979). Compare Siebel (2003).

[10]Searle (1969), Chapters 2.5 and 3.

involve the speaker properly *intending* that the illocutionary act be properly acted upon. A sincerity condition of an order, *e.g.*, is for the speaker to intend for the hearer to carry out the order. Finally, essential conditions are what distinguish one illocutionary act from all others: it's the defining feature of the act. For orders, the essential condition is that the speaker uses the utterance to try to get the hearer to do what the speaker wants.

For Searle, assertion has the following structure: "the preparatory conditions include the fact that the hearer must have some basis for supposing the asserted proposition is true, the sincerity condition is that [the speaker] must believe it to be true, and the essential condition has to do with the fact that the proposition is presented as representing an actual state of affairs" (*i.e.*, that it's true).[11] Furthermore, since the sincerity condition of asserting involves the psychological state of belief, for Searle, asserting counts as *expressing* the speaker's belief. Additionally, the preparatory conditions inform us what the speaker implies in asserting: namely, that the speaker *implies* that the hearer has reason for supposing the asserted proposition is true.

One serious problem with Searle's account is that he considers it a preparatory condition of assertion that the hearer has reason to suppose the proposition asserted is true. Suppose that the hearer knows, or is even *certain*, that some proposition is false. Suppose that the hearer is certain that pure water does not boil, at standard pressure, at 70°C.[12] If a speaker sincerely believes that water *does* boil at 70°C, then she may assert, "Water boils, at standard pressure, at 70°C." Since the hearer is certain that the assertion is false, the hearer

---

[11]Searle (1969, p. 64).

[12]It boils at 100°C at standard pressure.

cannot have any reason to believe that the proposition asserted is true. Thus the preparatory condition is not met. However, I think that it's clear that the speaker still asserts, though falsely. So it isn't a necessary condition for a non-defective assertion that the hearer will have reason to suppose that the asserted proposition is true.[13]

Other problems might arise when we characterize assertion as necessarily expressing a speaker's belief. Lies and bullshit (in the strict-ish sense of the term introduced by Harry Frankfurt) are both assertions, where the speaker puts forward a proposition as true, without expressing a speaker's belief in the proposition asserted.[14] If I intend to give false directions to the bus station, I may assert "Go North on 5^{th} Street," even though I don't believe that those are the correct directions. And although I'm lying, my utterance is still an *assertion*: lying is not a separate speech act from asserting.[15] While assertions may typically serve to express a speaker's belief, we should thus be cautious defining assertion as a speech act which expresses speaker beliefs.

A possible refinement to the idea that assertion expresses a speaker's belief in the proposition asserted is that the speaker must *intend* to assert something true. Lies and bullshit, then, would not count as properly assertoric because the speaker does not intend to assert something true in producing such utterances. After all, it's an intuitive feature of assertion that a speaker must at least offer the *impression* of sincerity. If a speaker clearly appears

---

[13]One might think that the speaker's asserting the proposition *necessarily* provides the hearer with *some* reason for supposing the proposition true. Perhaps, for example, it gives the hearer some reason to doubt her certainty that the proposition is false (and thereby provides some reason to suppose that it's true). However, if the hearer really is certain, then the mere act of hearing the assertion will not provide her with any reason to suppose the (false) proposition asserted is true.

[14]Frankfurt (1988).

[15]Though, I admit, one *could* try to argue that they're separate speech acts. However, I will not consider that possibility.

as asserting something that she doesn't believe to be true, the assertion will likely fail to be convincing.[16] Some evidence for this comes from noticing that hearers dismiss obviously insincere statements, something which is borne out in the challenge, "Come on, you don't really believe that." However, it doesn't follow that even obviously insincere assertions fail to be *assertions*. Rather, as Michael Dummett argues, "assertions are distinguished from other utterances in all being governed by the convention that we should try to utter only those whose descriptive content holds good."[17] That is, assertion is characterized by the convention where speakers should always intend to assert something true. On this sort of picture, then, one can argue that assertion *is* the speech act which expresses a speaker's belief, since it's difficult to see how speakers could properly intend to assert something true which they don't believe.

Dummett rejects this sort of characterization of assertion, though. He argues that assertion cannot be defined as a speech act that expresses a particular mental state (whether it's belief, judgment, knowledge, or something else), nor can it be defined simply in terms of a speaker's intention. Simply put, Dummett argues that a speaker's intention alone cannot determine which speech act is being performed by a given linguistic act: it's simply unable to bear the weight required to make the requisite distinctions. There are two problems, as he sees it.

First, a speaker's actual intention is irrelevant to which speech act is performed by a given linguistic act. Lies *are* assertions, after all. Even if we grant that good assertions require

---

[16]That is, such assertions don't fail to be assertoric, but there will be, *e.g.*, a failure of uptake or belief-adoption on the part of the hearer.

[17]Dummett (1981, p. 356).

certain preparatory conditions, such as the speaker at least giving the impression of sincerity, this is distinct from the speech act being an assertion: asserting and warrantedly asserting are two different things. Consequently, a characteristic feature of assertions, for Dummett, is that they can be used to deceive.[18]

The second problem is that a speaker may assert with many simultaneous intentions, none of which is necessarily the intention to utter something true. I may assert, for example, that broccoli is delicious to a child because I want the child to eat his vegetables. I might happen to believe that broccoli is delicious and so simultaneously intend to assert something true, because, perhaps, I prefer not to lie to my child in making such statements. However, the primary intention is simply to say something that will get the child to eat the broccoli; that I take the proposition as true is secondary. We can imagine similar cases where I assert to the child that broccoli is delicious when I believe it to be false. I may have the intention to convince the child to eat the broccoli (because I consider it healthy) by asserting something I take to be false, while simultaneously intending to offer the assertion as a demonstration to my spouse that one can assert something false and yet convince the child to do what I want. In both cases I've made an assertion, but in the latter I did not intend to assert something true.

Dummett thus rejects characterizing assertion strictly in terms of a speaker's intention and as expressing a speaker's belief (or some other mental state). Neither of these features can be necessary conditions since there are many obvious counter examples including lies and

---

[18]We should note that a number of speech can be used to deceive. But it's a characteristic feature of assertions that the dominant means of deceiving is through asserting rather than, say, commanding or requesting. Dummett notices this as well. Dummett (1981, p. 356).

even bald-faced lies.[19] However, he seems to think that it is a characteristic requirement of assertion that the speaker at least give the impression that she intends to assert something true.[20] In one sense this is uncontroversially true. When it's clear that someone is telling a story or acting on a stage, even though a number of linguistic acts appear to be assertions, competent hearers won't mistake the speaker for asserting. When an actor says, "There's a fire!" it's understood that he's not asserting that there's a fire in the theatre.[21] But it's problematic to characterize assertion as a speech act that necessarily requires speakers to at least give the impression of intending to say something true, since this seems to conflate making a normatively *proper* assertion and making an assertion. The former is sufficient but not necessary for the latter. One can argue that lies are *improper* assertions without being committed to treating lies as non-assertoric. Some obviously insincere assertions are still recognized as assertions, after all.

To recapitulate, Dummett takes the view that asserting requires the speaker to at least offer the impression of intending to assert something true because he considers some obviously insincere assertions as self-defeating. The issue here is properly understanding "self-defeating," but I will defer that discussion until Section 2.5. Behind Dummett's rejection of the speaker intention/belief expression view of assertion is that it ignores assertion as a fundamentally *conventional* practice. "Rather, there has to be some feature of the sentences which signify their being uttered with assertoric force; and an account of the significance of

---

[19]*E.g.* Sorensen (2007).

[20]Dummett (1981, pp. 298 - 300).

[21]We see this borne out in the incredulity an audience will have for an actor sincerely trying to warn the audience of a fire.

this feature – and thus the activity of making assertions – is to be given by describing the convention under which sentences possessing this feature are used."[22] What distinguishes assertion from other speech acts is that assertion is governed by the *convention* that one should utter sentences with the intention of uttering only true ones.[23] So how is this sufficiently different from the view Dummett rejects? First, although there's a convention that speakers, in making assertions, intend to assert only if the descriptive content is true, it isn't the case that one asserts only if the intention is present. Dummett describes how speakers typically learn the conventions. "When we first learn language, we are taught to make assertions only in the most favoured case, namely in that situation in which the speaker can recognize the statement as being true."[24] Children thus learn to assert only when their evidence conclusively establishes the truth of the assertion: assert that the mail has arrived only when one *sees* the mail carrier drop off the mail.

As we become more adept, however, we learn that assertions are acceptable in a wide variety of circumstances where the truth of the assertion is not conclusively established. Dummett's principal example is that assertions in the future tense – "I'll be at the party at 9," for example – "are never uttered in the situation which conclusively establishes their truth; and others, which we originally learned to utter only in such situations, we later learn to utter in circumstances in which we may turn out to have been mistaken."[25] In cases of assertions in the future tense, it's impossible to be in a position which conclusively establishes

---

[22]Dummett (1981, p. 354).
[23]Dummett (1981, p. 354 and elsewhere).
[24]Dummett (1981, p. 355).
[25]Dummett (1981, p. 355).

the truth of the proposition asserted.[26]

But that's not a problem. We must distinguish between the conditions that justify making an assertion and the truth conditions for what's asserted. Although the general convention is to intend to assert only true statements, truth isn't required for properly asserting in all circumstances, as I'll argue in Chapters 4 and 8. However, while being in a position to conclusively establish the truth of one's assertion isn't necessary, it's important to note that "we do not of course learn to make statements on no basis whatever, and, if we did, such utterances would not constitute assertions (even though they were made with the intention of uttering only true statements), because there would not be such a thing as acting on such statements."[27] Understanding the practice of assertion in terms of understanding what it means to *act* on assertions is key to Dummett's view.

Assertion, for Dummett, characteristically involves the speaker and the hearer taking on certain commitments. He argues that "[l]earning to use a statement of a given form involves . . . learning two things: the conditions under which one is justified in making the statement; and what constitutes acceptance of it, *i.e.* the consequences of accepting it. Here 'consequences' must be taken to include both the inferential powers of the statement and anything that counts as acting on the truth of the statement."[28] What it means to understand and accept an assertion as true is understanding what it is to act on the assertion. We test a child's understanding of "The mail has arrived" by observing their ability to act on the

---

[26]Of course, on some views of time such as eternalism, and given deterministic views of causation, it could be possible to be in a position to conclusively know the truth of a future tense statement.

[27]Dummett (1981, p. 355).

[28]Dummett (1981, p. 453).

assertion by, for example, fetching the mail. Moreover, what it means for a one to accept an assertion of "Your shoes are untied" is to bend down and tie one's shoes.

By making an assertion, one of the commitments taken on by the speaker is properly reacting to the assertion being corrected or challenged. In a sense the speaker commits herself to ensuring the truth of what they assert. Dummett argues, for example, that if a speaker's assertion is shown to be false, the speaker should be prepared to modify her statements. Sometimes a speaker will react by making her words fit the facts of the world by, for example, correcting an earlier assertion that a ball is blue by asserting that it is red if the ball really is red. But in other cases of having an assertion corrected, rather than withdrawing or correcting the statement, the speaker will simply change her behaviour to make his actions (the facts of the world) fit her words.[29] For example,

> "[i]f a man says, 'I am driving to Cambridge', and then learns that he is going in the opposite direction, his reaction may be not to withdraw what he said, but to make it true by turning round and going the right way: this does not make what he said any less of an assertion. What is essential, however, is not only that he should in fact be prepared, by altering either what he says or what happens to make what he says true, but that his intention in making his original utterance should be such that he regards it as mandatory thus to strive to make what he says agree with what is the case; no further choice remains to be made."[30]

Dummett's view of assertion, in a nutshell, is the following. It's a speech act governed

---

[29]One can see how Dummett's view differs from Searle's characterization of assertion as having a word-to-world direction of fit. See Searle (1979).

[30]Dummett (1981, p. 299).

by the general convention to intend to assert only true sentences. And although we learn to assert only in circumstances that conclusively demonstrate the truth of what's asserted, there are many circumstances where we properly assert short of having conclusive grounds for the truth of what's asserted. Finally, the making of an assertion creates commitments in both the speaker and hearer. The speaker takes on a commitment to ensure the truth of what's asserted, or else to respond accordingly when challenged (by perhaps giving reasons or evidence for its truth) or corrected (by perhaps retracting her statement). Meanwhile, the hearer takes on a commitment, when the assertion is accepted as true, to appropriately act on the new information if conditions require such action.

2.4   Brandom on Assertion

Robert Brandom's view of assertion is closely related to Dummett's view, and can be best approached by considering a question Brandom poses: "What is it that we're *doing* when we assert, claim, or declare something?"[31] Assertion is characteristically an action, a doing. It's not just that assertion is an action, though, since that's obvious: it's that assertion is to be understood in terms of its performative properties, and not (in the first instance) in terms of its associated semantic objects, like propositions that are conveyed via assertions, or the syntactic properties of sentences. On this view, we understand assertion as a performance, a concept I will deploy in Chapters 4 and 8 as a useful way of understanding the potential norms of assertion. What sort of doing is asserting? Assertion is, at its core, a social practice.

---

[31]Brandom (1983, p. 637).

And so understanding what it is to make an assertion must be understood against a socially-instituted backdrop of a dual structure of authority and responsibility. In asserting some proposition a speaker takes on responsibilities and authorizes hearers to take the assertion and *do* things with it. I will take these in turn.

When one asserts, one takes on specific socially-instituted responsibilities to justify one's assertions if challenged. This can include giving one's reasons or evidence for an assertion, or the use of further assertions which themselves may require justification. How far one must go in order to justify a claim is not determined *a priori*; instead, it's determined, in part, by what one's linguistic community requires for satisfying the justificatory requirements of defending an assertion in a given context. As Brandom puts it, what counts as adequate justification, on this view, is whatever the members of one's linguistic community "will let assertors get away with."[32] In low stakes situations, for example, very little will be required. If two people, Alice and Bob, arrive at a restaurant and Bob prefers hamburgers, he may ask Alice, "Do they serve hamburgers here?" Suppose that Alice asserts, "They do" and Bob challenges her assertion with, "How do you know?" If it's not terribly important for Bob that the restaurant serves hamburgers – if they don't, he will simply choose his next most preferred option given what the restaurant serves – then Alice's justification, "I saw it on the menu outside the front door" will likely suffice. But if the stakes are higher and Bob will eat at the restaurant only if it serves hamburgers, then Bob may find Alice's response inadequate: he may require further justification, such as corroboration from a restaurant

---

[32]Brandom (1983, p. 644).

staff member.[33]

It should be immediately apparent that if the vindication of assertions generally requires further assertions as evidence, then this suggests a potential vicious regress. However, Brandom thinks that this is not the case for three reasons. First, there may be a terminus for justification where an assertion offered as justification does not itself require justification. As noted previously, this is determined by one's linguistic community. If, for example, it is generally accepted in science that one need not defend the scientific method, then if one uses the method to justify an assertion, the justificatory requirements would be satisfied, terminating by appeal to the scientific method. Second, the demands for justification must themselves be reasonable and may stand in need of justification. The child-like "Why?" to any further reason in support of an assertion is patently unreasonable in many contexts. Third, we may recognize a special class of assertions which Brandom calls "bare assertions" that do not require justification. This is a general feature of, for example, pronouncements performed by people with special socially-conferred authority. A few examples include: judges making assertions of guilt, teachers making assertions of grades, baseball umpires calling a runner 'out,' and the Pope making assertions about matters divine. Brandom argues that such assertions do not stand in need of further justification beyond mere appeal to the relevant authority.[34]

The requirement for speakers to provide justification, in form of a responsibility, for assertions is central to Brandom's view. An important feature of this responsibility is that a

---

[33]Because, we may suppose, Bob may be so skeptical not even to trust Alice's showing him the menu she saw as they entered, but Bob will trust the say-so of an employee.

[34]Brandom (1983).

speaker's responsibility to vindicate an assertion is *conditional* on actually being challenged. Agents take on a kind of dispositional requirement to respond to challenges. And while an asserter immediately takes on the commitment to respond to challenges, whether an assertion is normatively appropriate will partly depend on how well the asserter would respond, were she to be challenged. I take up this important topic in more detail in Chapter 3.

In a nutshell, asserting for Brandom involves a dual structure of taking on certain commitments and authorizing one's hearer to take the assertion and do things with it. Most notably, the hearer is licensed to re-assert to others, or to use the assertion as a premise in forming inferences, which often leads to further assertions. These further assertions derive their appropriateness from the justification of the original assertion. Consequently, a hearer may take an assertion and defer "to the author of the original assertion the justificatory responsibility which would otherwise thereby be undertaken" by the new assertion.[35] Thus when a speaker makes an assertion which authorizes the hearer to take the assertion and use it, in some instances, to make further assertions, the justificatory duty of these further assertions falls to the author of the original assertion which licensed the subsequent series of inferences and assertions. So if Alice asserts to Bob that the restaurant serves hamburgers, and Bob then re-asserts to Charlie that the restaurant serves hamburgers, Bob could discharge the justificatory duty of the assertion by deferring to Alice.[36] Alice would therefore be responsible for providing the justification for Bob's assertion to Charlie. Of course, as she could do if her assertion were directly challenged, she could disavow the claim, leaving Bob

---

[35] Brandom (1983, p. 642), emphasis removed from the original.
[36] "Epistemic buck-passing" as Goldberg calls it. Goldberg (2006).

to decide whether he has his own reasons to offer in support of the assertion. A speaker thus takes on a responsibility for all the future assertions that are the fruit of an original assertion.

This is a particularly distinctive feature of Brandom's view of assertion.[37] However, it is also a possible target of criticism. Although a more complete discussion of responding to challenges such as "How do you know?" will be left until Chapter 7, that Bob can completely discharge the justificatory requirements of defending his assertion by passing the buck to Alice seems problematic. Cases like the restaurant example above are similar to instances of testimonial chains where one agent asserts to another, who re-asserts to another, who re-asserts to another, and so on. Suppose that the chain is five people long with the assertion originating with Alice, through Bob, Charlie, and terminating with Dan asserting to Erin. Suppose, then, that Erin doubts Dan's claim that the restaurant serves hamburgers and challenges the assertion with "How do you know?" According to Brandom's view, Dan could adequately respond to Erin's challenge by citing Charlie's assertion that the restaurant serves burgers. He would thus adequately respond with something along the lines of, "Charlie said so."

We should be skeptical of the adequacy of such responses. In some cases this is a plausibly adequate response: if we suppose certain background facts about the extreme reliability and epistemic care that Charlie takes with his assertions, then that Dan heard from Charlie may be a sufficient response to Erin's challenge. However, if it's mutually known that Charlie is a rather unreliable source of testimony, even if he happens to be felicitously re-asserting what he heard from Bob (and that it's true that the restaurant serves hamburgers), Erin

---

[37]I don't mean to suggest, though, that Brandom is the only proponent of this sort of view.

would find Dan's appeal to hearing from Charlie as bad, or at least insufficient, support for the assertion. In fact, this may give Erin positive reasons to *doubt* the quality of Dan's assertion. So there are important, but unspecified background facts required in order to make it plausible that a speaker can wholly discharge the justificatory duty of an assertion by appealing to the proximal or original author of the assertion. However, these facts are often absent in default contexts: speakers generally can't simply discharge their justificatory duty by merely appealing to their having heard it from someone else.

2.5   Taking Stock and Moving Forward

For my purposes, there are a number of features of Dummett's and Brandom's views that are useful in characterizing assertion. Assertion is an act, a doing, a *performance.* In Chapter 3 I will put this to work in arguing that the norms of assertion should respect that it is an act and the proper object of normative assessment, therefore, should be the agent performing the act. Asserting is putting forward a proposition as true, perhaps even as worthy of belief. Assertion is also, necessarily, social and constituted, at least in part, by social conventions. What makes an assertion appropriate will thus partly depend on context, including the social conventions of making assertions in the context. It's also important to notice that we use assertions to perform further actions including using assertions as premises in making inferences, particularly in practical decision-making. So that we do things with assertions ought to figure into properly characterizing the act. Furthermore, I agree with Dummett that assertion is not, at its core, the expression of a mental state, specifically belief. Although

asserting often expresses a speaker's belief, it's not a necessary feature of the act. I also agree that assertion involves the speaker taking on certain commitments to the truth of what's asserted, including, in part, a commitment to defend the assertion. And, moreover, asserting gives authority to hearers to take the assertion and do things with it, such as re-asserting it and using it as a premise in forming inferences.

I think that there are, however, features of both Dummett's and Brandom's views that we should try to avoid if possible. I disagree with Dummett that asserting what one doesn't believe is "self-defeating" in a significant sense. It depends heavily on what he means by "self-defeating." It's a conventional feature of assertion that a hearer will dismiss an assertion if the speaker makes it clear that the speaker doesn't believe what's being asserted. In this sense of hearer uptake, obviously insincere assertions are self-defeating. However, something is still asserted. The assertion may not have been normatively appropriate, in that it was a bad assertion *qua* norms of assertion, but it's an assertion nonetheless. Furthermore, provided that the speaker adequately gives the impression of sincerity, or at least sufficiently avoids giving an impression of insincerity, belief is not a necessary condition for assertion. After all, Dummett argues that a fundamental characteristic of assertion is that assertions can be used to deceive. And that they have this power depends, at least in some cases, on the speaker being able to assert something they don't believe.

There's also a problem with Dummett's claim that assertion is characterized by the convention that one should try to assert only true sentences. While this is a feature of many, perhaps even most, assertions, it's unclear in Dummett whether one can warrantedly assert

something false when one has good reason to think that what's asserted is true. I will argue in Chapter 4 that there are a number of circumstances where false, but well justified, assertions are warranted. However, in Chapter 8, I will go on to argue that there are some contexts where a speaker may warrantedly assert a known falsehood. Such assertions conflict with the claim that assertion is characterized by the convention that one should try to assert only true sentences.[38]

The principal point of my disagreement with Brandom is his view that the justificatory duty of an assertion falls to the original author of an assertion. I raised the problem of long testimonial chains. It's implausible that the last asserter in a long chain can fully discharge her duty to defend the assertion by appealing to the epistemic authority of the person from whom she heard it: more information about such cases would be required. Assertoric warrant isn't conferred so easily. If Dan has good evidence that Charlie is a reliable source of information, then this, in addition to having heard it from that Charlie, would constitute adequate justification for Dan's asserting to Erin. Having merely heard it from Charlie is insufficient: one cannot completely pass the buck to the original author of an assertion.

Furthermore, while I agree with Brandom that assertion has a dual structure of commitment and authorization, I strongly disagree with the "downstream" evaluation of an assertion in the form of a subjunctive clause on how a speaker would respond were she challenged. For Brandom, a present assertion is appropriate provided that it meets certain present nor-

---

[38]Though one could possibly retreat to the position that assertion is *generally* characterized by the convention that one should try to assert only true sentences.

mative requirements (such as, for example, that it expresses knowledge) and provided that the speaker would appropriately respond to challenges were she challenged. The problem, then, is that the normative assessment of a present assertion depends in part on the speaker's ability, and willingness, to adequately defend the assertion *if challenged*.[39] This seems wrong, though. In Chapter 8 I will argue that the norms of assertion function such that assertions are evaluated only at the time of assertion. Moreover, an assertion is proper, or improper, based on present features of the assertion, in terms of the agent's performance of the act of assertion. We evaluate an assertion when it happens, not partly based on contingent features of possible challenges and responses. If we suppose, for the sake of argument, that a given assertion is correct in a context only if the speaker knows the proposition asserted, then the assertion of a known falsehood will be improper irrespective of being challenged. More importantly, though, if one's assertion expresses knowledge and it's pragmatically and conventionally appropriate for the context (*e.g.*, it wasn't rude or impractical), then the assertion strikes us as warranted even if the speaker doesn't make good on the counterfactual requirement to adequately respond to challenges. Therefore, the normative appropriateness of the assertion doesn't depend, even in part, on how the speaker would respond if challenged. While I think that there are normative features governing how a speaker *ought* to respond to challenges, I don't think that these play a role in the norms of assertion.

The view of assertion taking shape from these reflections, then, is one on which an asserter

---

[39]One might think that the locus of evaluation for an assertion is in the present, since Brandom's "downstream" requirements consist in a present disposition to adequately respond to challenges. It's unclear to me what position Brandom takes. If his requirement is specified as a present disposition, then my objection doesn't apply to his view.

puts forward some proposition as true, and possibly even as worthy of belief by the hearer. The speaker necessarily takes on conventional socially-instituted commitments to what's asserted, which are manifested in a duty to defend the assertion.[40] Finally, a speaker asserting that p does not express a speaker's mental state regarding p, even though speakers often use assertions to accomplish such tasks as expressing belief or knowledge that p.

Throughout the epistemology literature, however, there is a substantially different methodology for characterizing assertion. Timothy Williamson, for example, argues influentially that we understand assertion by understanding its norms. What distinguishes assertion from other speech acts, such as commanding and requesting, he argues, is that it is governed by the following constitutive norm.

One must: Only assert that p if one knows that p.

Whether this is a fruitful research methodology for understanding the nature of assertion is not a question that I will take up in any detail. However, since the principal topic of this dissertation is an attempt to articulate and understand the norms of assertion, the question of what the norms of assertion may have to tell us about its character will be an interesting question to bear in mind, and a topic I will return to in Chapter 9. In Chapter 3 I take up the question of what sorts of norms the norms of assertion might be. Then, in Chapter 4, I return to the question of what the norms of assertion, specifically the *Truth Norm of Assertion* (TNA), might tell us about the nature of assertion itself.

---

[40]What it takes to adequately defend an assertion will be discussed in detail in Chapter 7.

Chapter 3

Norms and Normativity

3.1 Introduction

Since the central concern of this dissertation is the furthering of our understanding of norms of assertion, it will be helpful to say something about what is meant by "norm". This is important since much of what will be said about arguments for and against various proposed norms will depend on what sense of "norm" is being used. It's important to understand the relationship between the nature of a norm and how one argues for a particular norm. This, I think, has received insufficient attention in the debate over norms of assertion. So this chapter serves two purposes. First, I will discuss a number of central options for how we can understand norms and, specifically, the norms of assertion. Second, I will discuss how choices between these options constrain arguments regarding norms of assertion and the implementation of important epistemic concepts such as certainty, knowledge, truth, and reasonable belief (among others).

I will begin by discussing some basic features of norms. For example, a norm is a standard for the assessment of actions and, as such, must be appropriately action-guiding.[1] This leads into a discussion of a number of potential options for characteristic features of different norms. Norms may be, for example, constitutive or non-constitutive, defeasible or indefeasible, restrictive or non-restrictive, and simple or complex. I then discuss Williamson's influential characterization of the Knowledge Norm (KNA) as a simple single constitutive norm unique to assertion, and his analogy of norms with rules of games. I will argue that norms are importantly different from rules of a game and that Williamson's characterization of KNA using this analogy produces problems. The differences between norms and rules of games is one of the key restrictions on how the discussion of norms of assertion should proceed. This will lead into subsequent chapters where I take up, in turn, the various principal proposals for the norms of assertion such as the Truth Norm (TNA), Knowledge Norm (KNA), and versions of the Reasonable Belief Norm (RBNA). It will also heavily inform my analyses of the various arguments and evidence proffered for the various norms, as well as how I build the case for my proposed norm, the Supportive Reasons Norm (SRNA).

---

[1]In the following section I explain how norms may be about things other than actions, such as states of affairs. However, I'll argue that the relevant conception of a norm for the norms of assertion discussion are norms pertaining to actions.

3.2 Some Characteristic Features of Norms *qua* Norms of Assertion

The purpose of this section is to make clear some important features of norms insofar as we're interested in norms of assertion. What exactly is meant by a *norm* of assertion? This is not an easy question, for "norm" is a heterogeneous term with applications in a number of fields. There are moral norms (one ought not lie), social norms (one should hold the door open for other people), and norms for games (one must not throw a pass in football from beyond the line of scrimmage). In a minimal sense, norms express the rightness or wrongness of actions or states of affairs. Norms ranging over states of affairs, sometimes called norms of *being*, involve whether a certain state of affairs is good or ought to obtain.[2] Norms of action, on the other hand, tell us what to do: whether it's good to perform some action. Since I have argued in Chapter 2 that we should understand assertion as an action, a doing, it is norms of action that are most relevant for understanding the norms of assertion. So the first important working assumptions about norms of assertion is that they are norms of actions rather than norms of being.

In one sense, we might understand norms as more-or-less regulated *patterns* of behaviour. It's a norm of walking into a building that people will hold the door open for those behind them. This practice isn't codified as a law, or as explicitly required (at least anymore) by social conventions. However, it is an easily observable and relatively stable pattern of behaviour. But since this is merely a descriptive sense of norm, and not prescriptive, it is

---

[2]For a good introductory discussion of norms and normativity, see Gluer & Wikforss (2010, Section 1.2).

not how we will understand the norms of assertion. Norms of assertion do not attempt to merely describe how people participate in the practice of assertion; rather, they attempt to prescribe how people ought to assert in order to engage in the practice well. So the second important feature of norms of assertion is their prescriptive character: they're about the rightness, value, or appropriateness of an action (*i.e.*, assertion).

One question looms: why should anyone care that there are norms for a particular practice such as assertion? To say that a norm is *in force* is for an agent to accept her actions as subject to the normative assessments issued by the norm.[3] While an agent may fail to satisfy a norm in performing some action, she will incur, and be subject to, the normative force of the norm. The force of a norm concerns the nature of the "ought" and that a norm can figure as a premise in an agent's practical deliberation: that a norm can serve as a *reason* for action.[4] Suppose that it's a norm of dining that one should not eat potato salad with one's hands. If I consider this norm in force, then in deciding how to eat some potato salad, that the norm requires that I not do so with my hands will serve as a premise in my deciding to use a fork rather than my hands. A norm being in force provides an agent with reasons to perform one action rather than another, with the aim of satisfying the norm. Insofar as it is widely accepted that ought implies can, a norm can only have force if what it requires of an agent's actions is possible for the agent.[5]

A particularly critical feature of norms, as I conceive of them, is that they are action-guiding. Consequently, norms must be able to figure as premises in an agent's practical

---

[3]See, for example, Gluer & Wikforss (2010).
[4]See, for example, Gluer & Pagin (1999) and Boghossian (2008).
[5]However, I leave open the possibility that ought does not imply can.

deliberations. If we assume that an agent wants to engage in a practice properly, in that she wants to satisfy the norm of the practice and she considers the norm to be in force, then the norm of the practice should provide at least *some* guidance for how to go about satisfying the norm. An example will help. Suppose that Pierre is a hockey player. He's on a breakaway about to take a shot or perform a deke on the goalie. One might think that the norm for hockey is simply to score.[6] The problem is that this isn't sufficiently action-guiding: it depicts a desired outcome but neither illuminates nor constrains the means of achieving it. How should he go about attempting to score? Should he shoot to the glove side, shoot to the stick side, or try to deke? Even if Pierre consciously admits that he's endeavouring to score, and that he wants to engage in the practice of playing hockey well, the imperative to score doesn't provide any guidance for how he should go about attempting to satisfy the imperative. Of course, it's minimally action-guiding in that it directs him to score, rather than for him to shoot the puck over the glass and out of play, but the imperative is relatively impotent. If we suppose that Pierre selects an extremely improbable shot over an obviously highly probable shot, but still manages to score, the imperative is satisfied. But this seems to offend our notion of what he ought to have done, *i.e.* take the high percentage shot. That "ought," and not the outcome, is that for which we should use the notion of a norm.

Contrast the desire or imperative to score with a norm such as, "Take the shot that you think will have the highest probability of scoring." Now Pierre can use the norm to determine how he thinks he can best go about satisfying the norm. Suppose that Pierre knows that

---

[6]We might call this the "Score!" norm. Compare: "Buy low, sell high!" But at the time of purchase or sale, what actual purchases or sales does this enjoin or forbid? As I will discuss below, and again in Chapters 4 and 8, "Score!" and "Buy low, sell high!" are implausible norms but plausible *goals* of the respective practices.

the goalie is weak on shots to the glove side. Pierre can now select a glove side shot because that's what the norm tells him to do: this shot is the one that Pierre judges has the highest probability of scoring.[7] The norm helps specify what actions constitute rational means for achieving the goal of a practice: taking the shot that is most likely to score is the rational way to go about achieving the goal of scoring. This feature of norms will be especially important in the chapters to come since, I will argue, many of the arguments for or against various norms require a conception of norms which is not appropriately action-guiding. By making this feature of norms explicit, I will argue that much of the linguistic data and intuitions supporting TNA and KNA actually support a reasons-based norm such as the SRNA.

Thus far I have been discussing norms of *practices*, such as hockey. At this point we can make an important distinction between practices and *mere activities*.[8] Michael Rescorla argues that, "Every practice is associated with 'internal' standards of normative assessment codified by norms dictating how to execute the practice correctly."[9] An agent may violate such norms, however, and still be engaging in the practice: "A devious gambler might repeatedly cheat during a poker game. He plays poker incorrectly, but he plays poker nonetheless."[10] Rescorla argues that the "connection with internal evaluative standards differentiates practices, such as games or assertion, from mere activities, such as jumping, bathing, or holding hands."[11] There are no internal evaluative standards to jumping: once one satisfies the nec-

---

[7]Naturally, Pierre will want to play a mixed strategy for selecting his shots. If he always shot at the glove side, the goalie could learn this and adapt.

[8]Rescorla (2009). Rescorla is responsible for one of the more complete and most explicit discussions of characteristics of norms in the norms of assertion literature.

[9]Rescorla (2009, p. 101).

[10]Rescorla (2009, p. 101).

[11]Rescorla (2009, p. 101).

essary conditions for engaging in the activity of jumping (both feet leave the ground at the same time), there are no standards for doing it correctly or incorrectly. Of course, if we're speaking about jumping with a particular goal, such as jumping over a high bar, as in the sport of High Jump, then there are internal evaluative standards, but that's because we're speaking of a practice rather than a mere activity.

For Rescorla, the normative character of practices and activities is what separates them. He writes,

> Any practice engenders a three-fold division between actions that do not count as engaging in the practice, actions that count as engaging in it *correctly*, and actions that count as engaging in it *incorrectly*. A mere activity engenders only a two-fold division between actions that count as engaging in the activity and actions that do not.[12]

Practices have norms *internal to the practice*. That is, norms of hockey are norms *of* hockey, rather than more general social or moral norms, or norms of sports in general. Practices may also be governed by norms external to the practice, such as general moral, social, or rational norms. By contrast, activities are only governed by external norms: they lack internal norms. So the fifth important feature of norms, *qua* norms of assertion, is that we're interested in whether there are any norms specific and internal to assertion (that is, whether it's a practice or merely an activity). The one question will thus be whether, which, and how any external norms might govern the practice, and how such norms interact with any internal norms.

---

[12]Rescorla (2009, p. 101).

In some sense, even activities seem to have standards of assessment internal to the activity. For Rescorla, these are not norms; rather, they're simply *requirements* of the activity. A norm describes what's required to engage in a practice correctly or incorrectly; a requirement describes what's required to engage in a practice or activity at all. That is, requirements individuate an action rather than characterize it according to its fit with standards for acts of that kind. Take the activity of jumping: if Jane fails to have both feet leave the ground simultaneously, we wouldn't say that she is jumping poorly; rather, we'd say that she isn't jumping at all.[13] We would say that she isn't even engaging in the activity. So we can distinguish between norms, which describe what's required for engaging in a practice correctly, and requirements, which describe what's required to count as engaging in an activity or practice. All practices have both norms and requirements, whereas activities only have requirements.

In this section I have discussed five important features of norms, at least as the concept of a norm is important for the project of understanding and articulating norms of assertion. Norms are not just patterns of regulated behaviour: they're importantly normative in that they tell us what we ought to do. Since assertion is an action, we are concerned with norms about actions rather than states of affairs. We are also concerned with norms that are in force, in that they can figure as premises in an agent's practical deliberations about assertion. That is, the norms of assertion can guide what an agent asserts, or at least ought to assert. Consequently, norms should be appropriately action-guiding: empty norms such as "Score!" or "Buy low, sell high!" do not adequately guide an agent's actions in attempting to satisfy

---

[13]Intuitions may differ, though. I think that we could appropriately say that Jane is jumping poorly if she's attempting to jump but failing. We would both say that she's failing to engage in the activity, and that she's doing a bad job of jumping (in that she's failing).

such norms. Rather, such imperatives function better as goals of a practice rather than norms. Norms are rational means of achieving the goals of a practice. And, finally, we have distinguished between norms and requirements, and practices and activities. Norms of a practice, such as norms of assertion, are internal to the practice and can be distinguished from requirements for engaging in the practice.

I will later return to a sixth important feature of norms: that they are characteristically related to the goals of a practice and are importantly distinct from the goals. However, that discussion is best left until Section 3.5. Before we can discuss that point, I next turn to a number of options in how we can understand the characteristic features that we might think that norms of assertion should have. The questions which will occupy us in Section 3.3 are what sort of norms the norms of assertion might be. That is, whether they're constitutive or non-constitutive, defeasible or indefeasible, simple or complex, and restrictive or non-restrictive. After this discussion, in Section 3.4 I will turn to critically discuss Timothy Williamson's influential view of the norms of assertion as a single simple (restrictive and, arguably, indefeasible) constitutive norm unique to assertion. In section 3.5 I will end by laying out how I think we should understand the characteristic features of the norms of assertion.

3.3   The Normative Character of Norms

To say that a norm should be prescriptive, action guiding, internal to a practice, and have the other features discussed in Section 3.2 is not yet to say anything about what sort of

specific character proposed norms of assertion have or ought to have. There are a number of different forms norms can take, each with very different features. This section discusses four dimensions, each with two options: constitutive and non-constitutive, defeasible and indefeasible, simple and complex, and restrictive and non-restrictive. This list is by no means exhaustive, but it does cover the principal options for norms, and particularly norms of assertion.

Some norms have a special, logical relationship to the practice over which they range.[14] When a practice logically depends on its norms, the norms are said to constitute the practice: they're constitutive norms. John Searle distinguishes between *regulative* and *constitutive* rules. "[R]egulative rules regulate antecedently or independently existing forms of behaviour...."[15] Many forms of social behaviour have regulative rules. Dinner etiquette, for example, may dictate that one should only use a salad fork for one's salad, but the practice of eating dinner exists antecedently and independently of the salad fork rule. That is, the practice does not depend in any way on the existence of the rule: we could suppose that the rule never existed and the practice of eating dinner would remain intact.

The game of chess, on the other hand, is a paradigm example of a practice that has constitutive rules. That the pieces must move as they do, and what constitutes winning or drawing, creates the game of chess. As Searle notes, "The rules of football or chess, for example, do not merely regulate playing football or chess, but as it were they create the very possibility

---

[14]Until otherwise indicated, I will use "rules" and "norms" interchangeably. I take up important differences in Section 3.4.

[15]Searle (1969, p. 33). Others have sometimes merely distinguished between constitutive and non-constituve rules. See Gluer & Wikforss (2010).

of playing such games."[16] And so while "[r]egulative rules regulate a pre-existing activity, an activity whose existence is logically independent of the rules," nevertheless, "[c]onstitutive rules constitute (and also regulate) an activity the existence of which is logically dependent on the rules."[17] The game of chess cannot exist independently of these specific rules. Furthermore, Williamson writes, "a rule will count as constitutive of [a practice] only if it is essential to that [practice]: necessarily, the rule governs every performance of the [practice]."[18] In general, the norms of assertion literature has sought the constitutive norms of assertion: the norms that create and individuate the practice of assertion.

While regulative norms may be said to be non-constitutive, there are other ways that norms can range over a practice (*i.e.*, be in force) and yet be non-constitutive.[19] If we assume that the practice of assertion has a single constitutive norm, such as KNA (assert that p only if one knows that p), we don't think that satisfying the norm is sufficient for warrantedly asserting. Some assertions might be impolite or imprudent, for example, even if the speaker knows the proposition asserted. Although I might know that my friend has terrible breath, it might be inappropriate for me to tell him so because it will severely damage his feelings. So there are plausibly some general social norms which govern assertion, just as there are general norms of politeness that govern many practices. Such norms are in force for the practice, but are non-constitutive of that practice: we do not say that politeness is a norm

---

[16]Searle (1969, p. 33).

[17]Searle (1969, p. 34).

[18]Williamson (2000, p. 239).

[19]According to Searle's taxonomy, regulative norms are, by definition, non-constitutive. However, there's another sense in which norms not specific to a practice still "create" the practice. For example, we may think that politeness isn't a norm *of* assertion, and so isn't constitutive, but since politeness is an integral part of proper assertion, this regulative norm does "create" the practice in a sense: assertion wouldn't be the practice that it is if politeness wasn't one of its governing norms.

of assertion, since the norm is not internal to the practice. Instead, we say that assertion is governed by some external norms such as general conventional norms of politeness, for example. That is, one could say to a friend, "You have bad breath," hurt feelings and all, without one's assertion thereby being a non-standard case; and people *in general* could say rude things without destabilizing the practice or underlying conditions of assertion. So one issue moving forward will be to take care in distinguishing between whether a norm is internal to a practice – that is, a norm *of* a practice – and whether a norm is in force but external to a practice. In Chapter 8 I argue that there are social and pragmatic norms internal to the practice of assertion and, consequently, whatever proposed norm should accommodate this.

A second option for the normative character of norms *qua* norms of assertion is between defeasible and indefeasible norms.[20] Suppose that KNA is the norm of assertion such that one should only assert that *p* if one knows that *p*. Consider a case where Jim and Stacey are on a train platform.[21] Stacey is about to leave town and Jim is seeing her off. Jim sees a train about to depart and isn't sure if it's Stacey's train. The urgency of the situation is such that missing the train would be very bad, and being wrong about whether it's actually Stacey's train is a small cost. So we may suppose that Jim knows that he doesn't strictly *know* that it's her train, but since he knows that it's probably her train he nonetheless shouts "That's your train, run!" The pragmatic features of the context seem to override the knowledge norm. This therefore seems to be a case where one may warrantedly assert something one does not know, and would thus constitute a counter-example to KNA.

---

[20]Kvanvig (2009, 2011) considers this a central point of contention. He argues for a defeasible norm of assertion.

[21]From Williamson (1996, 2000).

But not so fast, argues Williamson. "Such cases do not show that the knowledge rule is not the rule of assertion. They merely show that it can be overridden by other norms not specific to assertion."[22] When a norm can be overridden by other, often more general, norms or considerations, then the original norm is said to be defeasible. Norms not subject to the possibility of being overridden are thus indefeasible norms. In Section 3.4 I will discuss Williamson's characterization of KNA in more detail since, despite what he says about the train case, he argues for an indefeasible norm of assertion: although Jim is reasonable in making his assertion, and perhaps even blameless (*viz.* DeRose's secondary propriety[23]), Jim still breaks the rule of assertion since he asserted something he didn't know. So the second consideration is whether we should consider the norms of assertion, whether constitutive or non-constitutive, as defeasible or indefeasible. In Section 3.5 I will briefly return to the issue of the interaction between norms internal to a practice and possible external norms which might override the internal norms. However, a more detailed discussion will be left until Chapter 8, where I propose and discuss the Supportive Reasons Norm of Assertion (SRNA).

Another question regarding the character of norms of assertion is the difference between simple and complex norms. According to Williamson, a norm of a practice is simple if there is a single constitutive norm of the practice and all other norms are a combination of the constitutive norm and considerations not specific to the practice. An act that satisfies the principal norm, whatever derivative norms it violates, will count as an instance of correct action in the salient sense. If I know that my colleague has terrible breath, then my assertion,

---

[22]Williamson (2000, p. 256).

[23]DeRose (2002), and elsewhere. I take up the topic of primary and secondary propriety in detail in Chapter 4 (see page 97), Chapter 7, and again in Chapter 8.

"You have bad breath," while rude, satisfies the norm of assertion. So it's appropriate in one sense, the salient assertoric sense, but inappropriate in the broader scope of polite statements to colleagues.

We might instead think that a practice, such as assertion, is constituted by one or more complex rules. Consider DeRose's famous bank examples.[24] We may suppose that in the low stakes case, the assertion that the bank will be open on Saturday is warranted.[25] But in the high stakes case, because the consequences for being wrong are dire, the agent is not warranted. So we may suppose a complex norm of assertion which states: "When stakes are low, one should assert that p only if one reasonably believes that p; when stakes are high, one should assert that p only if one knows that p."[26] Alternatively, we could suppose a complex norm of assertion that directly incorporates features which Williamson considers not specific to the practices, such as politeness. So it's a live possibility that the norm of assertion is complex. Perhaps the norm is something like the following: "One should assert that p only if one knows that p and it is sufficiently socially appropriate to assert that p." The norms of assertion literature has focused mostly on simple norms, but in Chapter 8 I will argue for a complex norm, SRNA.[27] It's important to note that a norm may be either simple or complex and yet be either constitutive or non-constitutive, and defeasible or indefeasible.

The fourth distinction is between restrictive and non-restrictive norms. Restrictive norms,

---

[24]DeRose (1992, 2002).

[25]I set aside the question of whether the agent *knows* that the bank will be open, which is the question DeRose is most interested in.

[26]DeRose draws a different inference about how knowledge shifts with context: one knows in the lower stakes case but not in the higher stakes case. Since DeRose ascribes to KNA, he thinks that this explains the shift in assertibility with changing contexts.

[27]However, there are other proposals for complex norms, and even some for *no* norms of assertion.

on the one hand, seek to restrict what one may warrantedly assert at the time of asserting. Such norms tend to have the structure: "One may assert that p only if p is ϕ," where ϕ is usually an epistemic concept such as knowledge, truth, or reasonable belief. KNA thus restricts what one is warranted in asserting in that the locus of evaluation is whether one, at the time of asserting, satisfies the norm of assertion. If one knows, then one satisfies KNA; if not, then one doesn't satisfy KNA. Michael Rescorla argues that most philosophers who've advocated constitutive assertoric norms, such as KNA, endorse restrictive norms.[28]

Non-restrictive norms, on the other hand, do not seek to restrict what one may warrantedly assert.[29] Quite simply, one can say whatever one pleases: whether an assertion is warranted will depend on how well one, for example, defends one's assertion if challenged. Rescorla argues for a non-restrictive norm of assertion with the following form: If one asserts p, then if D(p), one must ϕ(p). D(p) might be, for example, a speaker being challenged for asserting p, and ϕ(p) might be the speaker responding to the challenge by providing appropriate reasons for p. Non-restrictive norms make it difficult to evaluate, at the time of assertion, whether an assertion is warranted: it will depend on a great many factors including the original speaker's willingness and ability to adequately respond to challenges. In Chapter 8 I advocate a norm that is restrictive, but where a speaker incurs commitments to appropriately respond to challenges. However, the speaker's incurring these commitments is simply a feature of assertion and doesn't need to be incorporated into the norm in the way that Rescorla's non-restrictive norm does.

---

[28]Rescorla (2009).
[29]From Rescorla (2009).

Let's take stock. In this section I have presented four ways of distinguishing key characteristics of norms *qua* norms of assertion. Each dimension is a distinct option, which indicates that there are many forms the norms of assertion might take. In the next section I will discuss Williamson's prominent formulation of KNA as the unique single simple constitutive norm unique to assertion, what he calls a "C-rule" norm. He makes a close analogy with the norm of assertion with rules of games, and this analogy strongly influences the sorts of arguments and data he marshals for knowledge being the norm of assertion. However, as we will see, this produces a number of problems.

3.4   Williamson's C-Rule Norm of Assertion

Williamson sets out to answer the question, "What are the rules of [the game of] assertion?"[30] There's a close analogy, he thinks, between norms and rules of games. For example, violating a norm is like cheating at a game. He proposes that an "attractively simple suggestion" is that there is only one rule of assertion, a constitutive rule he calls the C Rule:

C Rule: One must: assert *p* only if *p* has C.

He argues that "knowledge" should stand in place of C: assert that *p* only if one knows that *p*. This norm thus normatively forbids asserting that *p* when one doesn't know that *p*. Although Jim's assertion for Stacey to run for a particular train seems appropriate, all things considered (including the pragmatics of the situation), it still violates the norm of assertion, strictly speaking, because Jim asserted when he didn't know. And although the C

---

[30]The bulk of the following discussion is from Williamson (2000, pp. 238 - 243).

Rule provides a necessary condition for properly asserting, knowing that p is not sufficient for warrantedly asserting that p. We've already considered cases where one may know that one's boss has bad breath, but it would be inappropriate to assert as much. Furthermore, constitutive rules don't provide the necessary and sufficient conditions for performing some act. For example, one may assert what one doesn't know without failing to assert: one just fails to assert properly. So we must distinguish the norms of a practice from its requirements, a distinction raised in Section 3.2 above.

Williamson thus argues for knowledge as the constitutive norm of assertion (KNA): One must: assert that p only if one knows that p. Since he considers it a constitutive rule, from our discussion in Section 3.2, the rule creates the practice of assertion: the practice would not exist without this rule. However, he also thinks that KNA is individuating. Assertion is the unique speech act whose unique rule is KNA. So there's a single constitutive norm of assertion, and it's the only speech act that has KNA as its (unique) norm. Below I will return to whether this is correct. But, furthermore, he also thinks that KNA is a simple norm of assertion: although assertion may be governed by many norms, "[a]ll other norms for assertion are the joint outcome of [KNA] and considerations not specific to assertion."[31] Moreover, despite some appearances to the contrary, he advocates an indefeasible norm.[32] Although there are some circumstances that make it reasonable to assert something one doesn't know, one still fails to satisfy the norm of assertion in the salient sense. Finally, since the locus of the normative evaluation of an assertion is at the moment of assertion, Williamson advocates a

---

[31]Williamson (2000, p. 241).

[32]Kvanvig (2011) also notes that although Williamson claims to support a defeasible norm, he seems to actually support an indefeasible version of KNA *qua* norms *of* assertion.

restrictive norm. Therefore, Williamson's KNA is a single constitutive, indefeasible, simple, restrictive norm of assertion. This will serve as a foil for my own views on the norms of assertion throughout this dissertation.

Earlier I promised to return to the question whether Williamson is correct to characterize assertion as the speech act uniquely and constitutively governed by KNA. I think that there is at least one other speech act constitutively governed by KNA. It's a rule of legal procedure that in a trial, lawyers questioning a witness may not make assertions: they may only ask questions. However, it's a norm of being a good lawyer that one ought not to ask a question to which one doesn't know the answer. The idea is that in the discovery and deposition stage of pre-trial proceedings, lawyers learn what they need in order to attempt to win the trial. At the trial stage during witness testimony, lawyers attempt to present the best case they can, and this involves not having witnesses surprise them with unanticipated information. So lawyers carefully choose their questions based on how they know the witness will respond to questions. Call this speech act a "lawyer question." It's plausible to suppose that the constitutive norm of this speech act is a knowledge norm: "Ask a lawyer question only if you know what the witness will answer."

Now recall that, on Williamson's view, one of the defining features of assertion is that it's the speech act uniquely governed by the Knowledge Norm: "$\phi(p)$ only if you know p." If lawyer questioning is a distinct speech act, and has the normative features that I have described, then the speech act is also governed by the Knowledge Norm. One should ask the question, "Did you see the defendant flee the crime scene" only if one knows that the

witness will answer a certain way (such as, "Yes."). If I'm correct that lawyer-questioning is a distinct speech act, and is constitutively governed by the Knowledge Norm, then it appears that Williamson would be wrong in claiming that assertion is the only speech act governed by the Knowledge Norm. However, I don't offer the example as a clear, knock-down objection to Williamson's claim. Rather, the point is merely to bring into question whether assertion really is the *only* speech act governed by the Knowledge Norm. In fact, even if it isn't, that's fine: it doesn't follow that assertion isn't governed by KNA, it just follows that KNA is not unique to assertion. And although KNA may not be the norm *unique* to assertion, it still may be the unique *constitutive* norm of assertion.

We may also question Williamson's characterization of the norms of assertion by analogy with rules of games. Williamson writes that, "one might suppose ... that someone who knowingly asserts a falsehood has thereby broken a rule of assertion, much as if he had broken a rule of a game; he has cheated. On this view, the speech act, like a game and unlike the act of jumping, is constituted by rules."[33] There are certain standardized patterns of acceptable behaviour in games, which are codified in the rules of a game. Often we find an official rule book: the rules of chess, for example. The rules make explicit which actions are permissible or prohibited, and which others are obligatory. So the rules make explicit which moves are correct. However, one can break some rules and still be said to be playing the game, as discussed above. So some rules may be broken while still engaging in the game. Williamson thinks that this is the case for asserting something one does not know: one is still asserting, albeit not with full propriety.

---

[33]Williamson (2000, p. 238).

This is vague territory, though. In some cases breaking a rule of a game means that one is not playing that game at all; in other cases, breaking a rule simply means that one is not playing that game correctly. In golf, in order to have a legal score in stroke play, one must hole-out (that is, appropriately hit one's ball into the cup on each hole) before moving onto the next hole. However, in many sub-communities, it is a common practice not to hole-out very short putts (such as those less than two feet). Golfers instead just pick up their ball and assume that they would have made the putt. This saves time, and does not generally alter the outcome, since the vast majority of such putts would have been made if attempted. But holing out is a central rule of playing stroke-play golf. So are such groups not playing the game of golf when they choose not to hole-out short putts? The intuition should be that such players are playing the game of golf, but perhaps not "fully" according to the rules.[34] Naturally, some rules will be more central to whether one is playing a game, rather than simply not playing a game correctly, than other rules, and it may not be clear *a priori* how to make these distinctions.

There are problems with the analogy between norms and rules of games.[35] Although the rule-book rules of chess define and create the game, and are thus constitutive rules, they simply describe the permissible moves of chess. They don't always have the sort of normative character we expect from norms. First, although the rules define what is required for winning in a technical sense (*i.e.*, check-mate), they do not describe how to best go about check-

---

[34]Alternatively, one could take a broadly Wittgensteinian position and say that the group is not, strictly speaking, playing "that" (*i.e.*, the official) game, but some other very similar game instead. I prefer the former view.

[35]See Maitra (2011) for a further discussion of problems with this analogy.

mating one's opponent. That is, they are not always appropriately action-guiding. Second, the constitutive rules of chess thus do not determine, beyond which moves are permissible, what will be good moves if one has the goal of becoming a grandmaster, for example. So although there are constitutive rules of chess, in that any permissible move of chess must adhere to these rules, the constitutive rules do not necessarily provide guidance for how one ought to move if one has further goals such as winning or becoming a grandmaster. That is, such rules do not necessarily describe how to play the game well: they merely describe how to play the game. Athletes can't simply read the rules of a game and know how to play the game well. That would require an understanding of topics such as proper strategy, which is not contained in the rules of a game. Norms, on the other hand, attempt to describe how to engage in a practice well, and not merely how to engage in the practice. And so the analogy of norms with rules of a game can only go so far. The constitutive rules of a game often lack the sort of normativity that we're interested in when attempting to understand potential norms of a practice, which includes games.

3.5   The Relationship Between Norms and the Goals of a Practice

So far I have discussed a number of important features of norms and how these features should restrict our discussion of the norms of assertion. In Section 3.4 I put some of these to work in criticizing Williamson's conception of the norms of assertion as analogous with rules of a

game. The problem, as I further argue below, is that rules of a game don't have the action-guiding and prescriptive normative character we want from norms. In this section I return to the fundamental distinction between norms and the goals of a practice and, moreover, I argue that norms derive their content from providing prescriptions for how agents ought to go about achieving the goals of a practice.

I argued earlier that we may view practices as partly constituted by their norms. What makes something an instance of *exercising* rather than as an instance of *competitive weight lifting* is partly a matter of what actions one ought to perform in context. Practices are also partly constituted by their goals: what one is trying to achieve within the practice. Exercising has the goal of improving health and well-being (among others), whereas competitive weight lifting has the goal of winning the competition and, plausibly, not the goal of improving health. Exercising generally doesn't have a goal of winning.[36]

There's an important connection between the goals and norms of a practice, I argue. Norms have a dual role in practices (apart from their constitutive role): on the one hand, they serve as standards of assessment for actions; on the other hand, they serve to provide a rational plan for achieving the constitutive goals of the practice. Norms prescribe good ways of taking action in attempting to achieve goals: following a norm is a good way to achieve a goal. The content of a norm flows from, in part, the (constitutive) goals of the practice and from what constitutes rational plans for achieving those goals. Norms disconnected from the goals of the practice simply don't make sense. It's not a norm of hockey that one

---

[36]Often competitive lifting results in a reduction in health: *i.e.*, injury! Of course, the training phase may improve health, but I mean to discuss the activity of competition.

should always score while skating on one foot because this doesn't provide a rational plan for achieving the goals of hockey (such as scoring). In fact, it probably hinders achieving the goals. Consequently, what one ought to do necessarily depends on what goal one is attempting to achieve. If the goal of hockey is to score, rather than shoot the puck at the referee, then the norm should be "Take the shot which you think will have the highest probability of scoring," rather than "Take the shot which you think will have the highest probability of hitting the referee."[37] If we change the goal of a practice, we change the norm. A norm is, I argue, what it means to *properly aim at the goals of the practice*. Taking the shot with the highest probability of scoring is properly aiming at – *i.e.*, a rational means for obtaining – the goal of the practice.

Norms must be separate from the goals, though. Let's suppose that the goal of believing is to have true beliefs.[38] Still, the norm for belief can't be, "One should only believe that p if p." That would be identifying the norm with the goal; and it would not be action-guiding in the appropriate sense. Goals don't outline rational plans for good ways of going about achieving them. If I know that check-mating my opponent is the goal of chess, the goal doesn't give me any idea how I should best go about doing so. That's what norms do. This is why reading a rule-book and learning the goals of a game won't inform an agent how to best go about achieving the goals of the game. Consequently, norms must be distinct from goals. So the norm for believing might be, "One should believe that p only if p is best supported by

---

[37]One might question whether the goal of hockey is scoring goals or, for example, winning games. For the sake of argument, assume that it's winning games, but scoring goals is an intermediate goal in service of the overarching goal of winning games.

[38]I set aside the issues of whether knowledge, or even certainty, is the real goal of belief (or if belief even has any goals). There's a rich literature on this topic.

one's evidence, e, and e is sufficiently strong." This outlines a rational plan for attempting to achieve the goal of belief. It also provides a standard for assessment: in believing that p, is p best supported by subject S's evidence? If yes, then it's a good belief; if no, then it's not. Epistemologists tend to call this property of belief *justification*. Justification is simply the name given to "good" beliefs, whether they're true or not.[39] One upshot of keeping norms and goals distinct is that we accommodate the sense that one can act unimpeachably while obtaining the wrong outcome, a distinction I will put to work in Chapters 4 and 8.

One might naturally think that a justified belief, *i.e.* one that conforms to the afore-mentioned norm, that fails to be true is still an inappropriate belief in a salient sense. It fell short of the goal, after all. This is an extremely important point, though: *achieving the goal is distinct from satisfying the norm*. One can rationally and competently carry out a plan to achieve a goal while still failing to obtain the goal. The act is still appropriate in the normative sense, even though it falls short of the goal. Perhaps the act is not fully praiseworthy, in that an act would have been better *ceteris paribus* had it also achieved the goal, but it doesn't follow that the act was therefore "bad" or normatively criticizable in any meaningful sense. This is a mistake that I will diagnose in TNA, KNA, and CNA positions throughout this dissertation, in various forms.

---

[39]I also set aside the many and varied debates over the nature of epistemic justification.

3.6 Conclusion

In this chapter I have attempted to clarify the difficult concept of a norm. I've argued that we should understand norms, *qua* norms of assertion, in a fairly specific sense. For example, norms of assertion are norms of action, rather than states of affairs, they must be in force, and they must also be appropriately action-guiding. In Section 3.5 I went further to say that norms must be action-guiding in the sense that they provide rational plans of action for going about achieving the goals of a practice. So, in a sense, not only does "ought imply can," but "ought implies *how*." Consequently, norms must also be able to serve as reasons in practical reasoning. The reason should now be clear: since norms provide prescriptions for how one should go about trying to achieve the goals of a practice, these prescriptions fit nicely as premises in practical reasoning.

I have also discussed a number of options for the character of the norms of assertion. Norms may be constitutive or non-constitutive, defeasible or indefeasible, simple or complex, and restrictive or non-restrictive. With these in place, I presented Williamson's prominent characterization of KNA as the unique single simple constitutive, and restrictive norm unique to assertion. I also discussed his analogy of the norms of assertion with rules of a game and gave reason to doubt the usefulness of this analogy: rules of a game are too different from norms of a practice. Rules don't tell us how to play a game well, but norms do.

I also argued that the content of a norm is derived, in part, from what it means to rationally attempt to achieve the goals of a practice. This will serve as an important methodology for moving forward. Advocates of various norms for assertion have noticed a growing body

of linguistic data points relevant to articulating norms of assertion. Williamson, for example, writes, "Much of the evidence for the knowledge account comes from the ordinary practice of assertion."[40] The upshot of the discussions of this chapter is that properly understanding the restrictions on norms *qua* norms of assertion will restrict the sorts of arguments and evidence brought to bear in the contemporary debates on the correct norms of assertion. For example, proposed norms of assertion must be appropriately distinct from the proposed goals of assertion. The methodology for articulating the norms of assertion has often been to notice the linguistic data first, and then seek a norm that best explains it. I have suggested a different method: by understanding the goals of assertion, we can understand the norms. However, I have not yet commented on plausible specifications of the goals of assertion. Attempting to articulate plausible goals for assertion will be a recurring topic for the remainder of this dissertation. A very plausible goal is that one aims to assert only true statements. In fact, it's also plausible that assertion aims at asserting known propositions. This will be a topic taken up in Chapter 4, where I argue against the Truth Norm (TNA) partly on the grounds that its advocates use arguments that rely on mistakenly equating the possible goals of assertion with its norms. I also return to discuss this issue in more detail in Chapter 8.

---

[40]Williamson (2000, p. 243).

Chapter 4

Truth and Assertion's Aims

4.1   Introduction

The principal goal of this chapter is to better understand the relationship between truth and norms of assertion. I will go about this by considering two prominent defences of the view that truth itself is the constitutive norm of assertion.

>   TNA: assert that p only if p is true.

It has been a consistent feature in debates over norms of assertion that asserting falsely is likened to cheating at the game of assertion.[1] This implies that assertion has a *factive* norm – one for which truth is a necessary condition. Crispin Wright and Matthew Weiner argue that the central epistemic concept to warranted assertibility is the truth of what's asserted.[2] Wright famously argues that truth's role in warranted assertibility carries implications for

---

[1]Williamson (2000).
[2]Wright (1992) and Weiner (2005).

theories of truth more generally, and in particular that the *deflationary* theory of truth cannot be correct. I'll argue, building on where I left off in Chapter 3, that although it's plausible that truth is the central *goal* of assertion, truth is distinct from the *norm.* Wright's argument depends on this conflation. Once this error is corrected, I argue, deflationism seems adequate to fill truth's role in norms of assertion, and so Wright's "inflationary argument" fails.[3]

Weiner, by contrast, argues for the truth norm more directly. He argues that the motivations for the Knowledge Norm (KNA) actually support TNA. The linchpin of this argument is the contention that KNA gets some cases wrong which TNA gets right. There is a residual problem for TNA, though: lucky guesses are true but unwarranted assertions. Weiner attempts to defuse this problem by appeal to DeRose's distinction between primary and secondary propriety of assertions (and actions more generally). I will argue that Weiner's solution is vulnerable to still more problematic cases of lucky guesses. A further problem looms on the horizon: the lottery paradox. KNA advocates hold up the intuitive unassertibility of the statement, "Your ticket did not win" made about a lottery ticket before the speaker has information on the results of the draw. Even if the assertion turns out to be true, and even if it isn't really a lucky guess, it still seems unwarranted. TNA may not be able to fully explain this.

---

[3]Even if one were to eventually reject a deflationary theory of truth, my argument implies that one shouldn't do so because of Wright's inflationary argument. Although I won't discuss the matter, I ultimately reject deflationism since it seems to rule out a distinction between knowing that p and knowing that p is true.

4.2 Connections Between Truth and Norms of Assertion

Crispin Wright has famously argued that deflationary theories of truth are unstable and have a tendency to "inflate" under pressure. The key to his inflationary argument is the claim that deflationists are committed to truth playing a particular role in norms of assertion. The problem is that "is true" and "is warrantedly assertible" coincide in normative force, but diverge in extension. Since the deflationist must appeal to something beyond the disquotational schema, deflationism is unstable. But, I argue, we must remember that there are a variety of forms of deflationism. While Wright's argument succeeds against certain forms of the view, it remains to be seen whether it succeeds against deflationism *tout court*. In this first part of the chapter, I argue that it does not: Wright's argument depends on attributing a view to deflationist which the deflationist is not committed to taking. Consequently, I argue that Wright's inflationary argument doesn't succeed against deflationism *tout court*.

There are myriad varieties of deflationism, mostly trading on small but important differences in what functions the truth predicate has. All forms, though, ascribe to some version of the Equivalence Schema (ES), sometimes referred to as the Disquotational Schema (DS):

ES: "P" is True iff P

Some have claimed, for example, that the ES means that asserting "That's true" when referring to the proposition *Chicago is large* means the same as asserting the proposition, "Chicago is large." But compare the following two exchanges between friends.

Fred: "Chicago is large."

Mary: "Chicago is large."

Fred: "Chicago is large."

Mary: "That's true."

If it's really the case that Mary's asserting, "That's true" (referring to Fred's assertion that Chicago is large) is equivalent in meaning to her asserting, "Chicago is large," then we should consider both of these exchanges as equivalent. One might think that we don't, at least not intuitively. However, the difference is one of pragmatics, not semantics. In the first exchange, Mary may have made her assertion without hearing Fred's, but in the second exchange, her assertion clearly acknowledges Fred's contribution. Now, while an adherent to ES need not be committed to treating the two acts of asserting as *meaning* the same, some versions of deflationism do. One version of deflationism holds that we use the truth predicate as a prosentence, known as a *prosentential theory of truth*. Much like a pronoun like *he* stands in for Fred, *that's true* stands in for the expression, "Chicago is large." However, although the semantic content of the sentence and the prosentence are the same, their pragmatic features differ in the way described above: the use of the prosentence importantly acknowledges the other speaker's contribution in a way that the use of the sentence doesn't. The use of the prosentence is just a syntactic shortcut for uttering the full sentence. Since this is the view Wright attacks, it is the view I will present as the deflationary view, recognizing that there are alternative formulations.

An important feature of the prosentential version of deflationism is that using the pros-

entence "That's true" is a claim about the semantic content of the antecedent sentence (*i.e.* that "Chicago is large" expresses an extralinguistic fact about the size of Chicago) rather than the antecedent utterance itself. As Dorothy Grover writes,[4]

> Many other truth theories assume that a sentence containing a truth predication, e.g., 'That is true,' is about its antecedent sentence ('Chicago is large') or an antecedent proposition. By contrast, the prosentential account is that 'That is true' does not say anything about its antecedent sentence (e.g., 'Chicago is large') but says something about an extralinguistic subject (e.g., Chicago).

The point is that Mary's use of the prosentence isn't saying anything about the appropriateness of *Fred's* utterance of "Chicago is large," but, rather, is saying something about the size of Chicago: that is, about the proposition expressed in Fred's utterance. Thus the prosentence is about the semantic content of the assertion, rather than about the assertion (*qua* act) itself.

Deflationism about truth, of the prosentential stripe, is roughly the claim that the ES captures everything about truth that is of interest to philosophical logic and philosophy of language. It claims that to assert that a proposition is true is merely to repeat the assertion of the proposition. "'Snow is white' is true" is equivalent in meaning to merely asserting "Snow is white." (Though, again, there may be pragmatic differences.) Furthermore, it is the denial that truth amounts to anything more than the ES insofar as truth is not "analyzable" beyond the ES. Deflationists sometimes claim, for example, that many of the considerations

---

[4]Grover (1992, p. 221).

that motivate theories of truth such as correspondence are claimed not to be the "job" of truth.[5] How true statements may connect with the world or be caused by features of the world is not something for which a theory of truth must account. Instead, it is suggested, this may be the purview of theories of language or metaphysics, but not of truth *per se.*

The ES is thus open to expression as a semantic claim about the equivalence in meaning between "'Snow is white' is true" and "Snow is white," with the obvious implication being that the predicate 'is true' is, in some fairly profound sense, empty. More precisely, the deflationist holds two theses, one positive and one negative. The positive thesis is that the ES "is, near enough, wholly explanatory of the predicate, 'true'[;] that, for one who understands a mentioned sentence, the claim that it is true comes to the same thing as its assertion."[6] That is, truth "is merely a device of 'disquotation'; rather than expressing any substantial characteristic of sentences, the effect of its predication, at the metalinguisitic level, is just to accomplish what, within the object language, we can accomplish by asserting the sentence in question."[7] As Quine puts it, one aspect of the truth predicate is its role as a device for semantic ascent from the object language to the metalanguage.[8] When we say, "'Snow is white' is true," we mean the same as if we had just removed the quotation marks around 'Snow is white' and merely said, "Snow is white." Deflationism is thus sometimes called the "disquotational" theory of truth.

To quote Wright, the negative deflationist thesis is "that truth is not a 'substantial prop-

_____

[5]Horwich (1998), for example.
[6]Wright (1992, p. 14).
[7]Wright (1992, p. 14).
[8]Quine (1986).

erty,' whatever that means, of sentences, thoughts, and so on, but merely a device for accomplishing at the metalinguisitic level what can be accomplished by an assertoric use of the mentioned sentence" in the object language.[9] Whether truth is really a substantial property, or not, will not be of primary interest for our purposes. The deflationist thus claims that the predication of "true" to a sentence is connected to assertion in that asserting that a sentence is true can also be accomplished by merely asserting the sentence. We can also notice by a cursory observation of how we use "true" in conversation that one central use of the truth predicate is for the endorsement of assertions. If my friend says, "It's windy today," I can agree by saying, "That's true." So truth seems, at the very least, to be a predicate suited to the appraisal or endorsement of assertions. This might suggest a connection between the concept of truth and norms of assertion, which is something that Crispin Wright uses to "inflate" deflationism.

4.3   Wright's Inflationary Argument Against Deflationism

When we combine the ES with the view that truth is a means for endorsing assertions, Wright argues, we end up committed to truth's having a particular normative role in assessing warranted assertibility. The heart of Wright's negative case is that deflationism adopts both antecedent views, but is inconsistent with the consequence they engender. Wright proposes to "characterize as a *practice* any form of intentional, purposeful activity, and as a *move* any

---

[9]Wright (1992, p. 15).

action performed within the practice, for its characteristic purposes."[10] He also distinguishes, rightly, between prescriptive and descriptive norms. "A characteristic of moves in a particular practice is a descriptive norm if, as a matter of fact, participants in the practice are positively guided in their selection of moves by whether a proposed move possess that characteristic."[11] Descriptive norms thus describe what norms people actually use in choosing their moves within a practice. If people actually assert based on whether they think that their assertion amounts to knowledge, then the knowledge norm (assert that p only if one knows that p)[12] is a descriptive norm of assertion. "By contrast, a characteristic of moves supplies a prescriptive norm just in case the reflection that a move has that characteristic provides a (defeasible) reason for making, or endorsing, or permitting it, even if such reasons tend, for the most part, to go unacknowledged by actual participants."[13] Implicit in the phrase "reason for making..." is the idea that it's a normatively good reason to assert if one thinks that one's assertion satisfies the norm. So if one thinks that one's assertion has the characteristic, which is the normative property of the move, such as truth, then this provides a defeasible reason for thinking that one's assertion is warranted (or appropriate, or justified).

To recapitulate, then, switching from talk of characteristics or properties, "a *predicate* F, is (positively) descriptively normative just in case participant's selection, endorsement and so on of a move is as a matter of fact guided by whether or not they judge that move is F – a judgment whose bare possibility is neutral, presumably, on the question whether or not the

---

[10]Wright (1992, p. 15), emphasis in the original.
[11]Wright (1992, p. 15).
[12]See, for example, Williamson (2000), DeRose (2002), Hawthorne (2004), and Turri (2011).
[13]Wright (1992, p. 15).

predicate expresses a substantial property. Likewise a predicate is prescriptively normative just in case the selection, or endorsement, of a move *ought* to be so guided within the practice concerned."[14] So Wright claims that the deflationist is committed to the position that the truth predicate is both positively prescriptively and descriptively normative of any assertoric practice.

He argues, though, that we can draw an even stronger conclusion for the deflationist's position: for the deflationist, "true" and "warrantedly assertible" coincide in positive normative force. Two predicates coincide in normative force if "each is normative within the practice and reason to suppose that either predicate characterizes a move is reason to suppose that the other characterizes it too."[15] That is, it follows from the ES that "any reason to think that a sentence is T may be transferred, across the biconditional, into reason to make or allow the assertoric move" to assert that p.[16] So Wright argues that the deflationist is committed to the thesis that both truth and warranted assertibility are norms of assertion and, moreover, coincide in normative force.

This produces a problem for the deflationist, since the deflationist must insist that "the only substantial norms operating in assertoric practice are norms of warranted assertibility, and that the truth predicate can indeed mark no independent norm."[17] To claim that truth is an independent norm from warranted assertibility would mean that the deflationist thinks of truth as a "substantial property," which is to deny the deflationist's own negative the-

---

[14]Wright (1992, p. 16), emphasis in the original.
[15]Wright (1992, p. 18).
[16]Wright (1992, p. 18).
[17]Wright (1992, p. 18).

sis discussed above. The problem arises from truth and warranted assertibility necessarily registering distinct norms of assertion: although they're normatively coincident, they are extensionally divergent.

Here's how Wright demonstrates the extensional divergence of "is true" and "is warrantedly assertible."[18] Since constituents of the ES must allow for significant negations, it follows that

    (i) "It is not the case that P" is T if and only if it's not the case that P.

Because of the biconditional in the ES, we can also conclude the following:

    (ii) It is not the case that P if and only if it is not the case that "P" is T.

And it follows from the transitivity of the biconditional,

    (iii) "It is not the case that P" is T if and only if it is not the case that "P" is T.

Here's the crux of the argument: "Reflect, however, that (iii) must fail, right-to-left, when 'is T' is read as 'is warrantedly assertible.'"[19] He goes on to write, "[h]ence, since (iii) holds good for the T-predicate, we have to acknowledge some sort of conceptual distinction between 'is T' and 'is warrantedly assertible.'"[20] And so the extensions of "is true" and "is warrantedly assertible" diverge.

The key to the inflationary argument is that what explains the normative coincidence but extensional divergence is some property of truth. Thus, since there is some property of truth

---

[18]Wright (1992, pp. 19 - 20).
[19]Wright (1992, p. 20).
[20]Wright (1992, p. 20).

that must explain this difference, and the deflationist is committed to truth not having any properties other than what is characterized by the ES, then the deflationist cannot account for this property of truth. The deflationist is committed to truth having this property and that, therefore, their position is inconsistent. And thus deflationism has a tendency to "inflate."

4.4   Analyzing Wright's Inflationary Argument

Insofar as it's appropriate to read "is true" as "is warrantedly assertible" in (iii), Wright successfully inflates the deflationist's position. And some deflationists seem to have made claims that would warrant this treatment.[21] Some deflationists have claimed that we use the truth predicate as a prosentence – the view on which "That's true," when uttered following an assertion, typically functions in the way that "Yes" functions when uttered following a question (of the appropriate sort). Hence, on the prosentential variety of deflationism, the expression "is true" is a means of *endorsing* assertions in the richest sense; that is, it's a syntactic shortcut for *reiterating* them. This seems tantamount to ascribing warranted assertibility to a sentence or proposition; so deflationism of this sort lends itself to the view that truth has a role in norms of assertion. Indeed, Wright holds some deflationists to have claimed that truth *is* a (indeed, *the*) norm of assertion. And this seems to get the deflationist in trouble. From the biconditional in ES, any reason to think that p is true is a reason to think that p is warrantedly assertible. So "is true" and "is warrantedly assertible" seem coincident in normative force. However, their extensions diverge: intuitively, lots of truths

---

[21]As noted, Wright seems to have Paul Horwich as the main target.

are not warrantedly assertible, owing to such factors as ignorance. The problem for the deflationist is that truth therefore must have a more substantial content than the norm of assertion alone, in order to explain the extensional divergence from warranted assertibility.

In this section I will present my principal criticism of Wright's argument – or, more accurately, my criticism of the attempt to generalize Wright's argument to deflationism across the board. Insofar as Wright properly characterizes some deflationists' position and their commitments, I agree with his inflationary argument. It strikes me as entirely effective against the position he wants to engage. My interest, however, lies in the bare question of whether truth itself must express a norm of assertion, and whether deflationism is really committed to truth being a norm of assertion. Hence I want to explore whether Wright's argument establishes this unconditionally.

Although some deflationists may have claimed that a deflationary theory of truth is committed to an intimate connection between truth and warranted assertibility, I'll argue that deflationists need not be so committed. At least, there's nothing in ES itself that commits a deflationist to the view that truth expresses a norm of assertion, or even that truth has any role in norms of assertion. But since the truth predicate does have a very important role in our linguistic practices, one might think that the deflationist must explain truth's role in terms of norms of assertion. I'll argue that this is mistaken. An adequate theory of truth will need to posit a concept of truth that is able to fill its role in our linguistic practices. If the deflationist is unable to do this, then this is a mark against the view. As I'll argue, I think that deflationary truth is sufficiently robust to meet this requirement.

The crucial point, though, is that the deflationist does not need to explain *why* truth has the particular role that it does in our linguistic practices, particularly in terms of norms of assertion: that's the job of theories of warranted assertibility. In service of this criticism, I will focus on what I take to be Wright's incorrectly imputing a view to deflationism concerning the use of the truth predicate as a prosentence in *endorsing* assertions.

Wright takes there to be two principal platitudes about the truth predicate and warranted assertibility: first, "asserting a proposition is claiming that it is true, and second that every proposition has a significant negation."[22] Both of these seem correct. On the one hand, it's a common definition of assertion as a speech act where one puts forth a proposition as true. Assertions are truth-apt: capable of being described as true for false. Both of these are analytic features of assertion. And on the other hand, having a significant negation is simply a feature of propositions. Both of these platitudes seem uncontroversial. However, Wright thinks that these two platitudes "ultimately ensure that truth and warranted assertibility coincide in normative force but are potentially extensionally distinct...."[23]

One problem arises from an ambiguity in what Wright means by a norm of a practice. He loosely defines norms in terms of whether moves in a practice are (descriptive norms) or ought to be (prescriptive norms) guided by whether an agent judges a move to have a characteristic. Wright proposes that the deflationist claims that both truth and warranted assertibility are prescriptive (and descriptive) norms of assertion. Two norms are coincident in normative force, according to Wright, if aiming at one is, necessarily, aiming at the other, and

---

[22]Wright (1992, p. 24).
[23]Wright (1992, p. 24).

that reason for judging that a move has one characteristic (*e.g.*, truth) provides a defeasible reason for judging that the move has the other characteristic (*e.g.*, warranted assertibility). But this is an exceptionally weak definition, one which masks the distinction between a norm and an aim (or *goal*) of a practice. Wright has claimed that "warranted assertibility" is a norm of assertion. But this is mistaken: warranted assertibility is the name we give to the norm, or constellation of norms, for assertion. "Warranted assertibility" isn't a norm: truth (as in: assert that p only if p) or knowledge (as in: assert that p only if one knows that p) are norms of warranted assertibility.

There's another important reason to distinguish between norms and goals; namely, that one can achieve the goal of a practice while violating the norm. One might think that it's analytic that achieving the goal of a practice provides a (defeasible) reason for endorsing the action, much like Wright argues that any reason to think that an assertion is true (the goal) is a reason to think that it's warrantedly assertible (the norm). But the extensional divergence of "is true" and "is warrantedly assertible" shows why this is mistaken. And, moreover, the case of lucky guesses makes it more obvious: one can truthfully assert while failing to meet the norm of assertion.

More importantly, a reason to think that an assertion is true is *not* thereby a (defeasible) reason to think that the assertion is warrantedly assertible: more information is required in order to make the latter judgment. The norms of assertion are more demanding. Although a reason to think that a proposition is true provides *some* reason to think that the proposition is warrantedly assertible, Wright's use of "reason to think" is ambiguous. For example, my

reasons for thinking that the emperor has no clothes might be epistemically unassailable, but these aren't reasons for thinking that it's warrantedly assertible to say so to the emperor (since he doesn't take kindly to such news). Certainly, conclusive reasons to think that a proposition is true are not therefore conclusive reasons to think that the proposition is warrantedly assertible: I may know with certainty that p is true and yet there are no contexts where it's warranted for me to assert that p. Admittedly, this is partly Wright's point in arguing for the extensional divergence of truth and warranted assertibility. My objection has to do with Wright's claim regarding the normative coincidence of truth and warranted assertibility. However, I'll have more to say below.

Now, if what it means for an assertion to be warranted is that the speaker properly aims at the assertion being true, then we should expect that aiming at truth is, necessarily, aiming at warranted assertibility.[24] This is simply analytic of the relationship between the norm and the aim of a practice. However, from this analyticity it does not clearly follow that "any reason to think that a sentence is T may be transferred, across the biconditional, into reason to make or allow the assertoric move" to assert the sentence.[25] Wright argues that this simply follows from three considerations: (i) the ES, (ii) the normative coincidence of truth and warranted assertibility, and (iii) that to say that some assertion is true is to endorse it.

How does this argument run? Begin with the ES.

(1) "p" is T iff p.

---

[24]By "properly" aiming I mean something like, "doing well in aiming at the goal." One can aim at an archery target by closing one's eyes, spinning around in a circle, then letting the arrow fly. But this wouldn't be "properly" aiming at the target. One could say of assertion, for example, that following one's evidence is properly aiming at truth, or one could say that one only properly aims at truth when one successfully obtains truth. There's room for debate here.

[25]Wright (1992, p. 18).

Suppose we grant for the sake of argument that the operator WA has the appropriate logical properties to allow its introduction on both sides of an equivalence. Then the deflationist is committed to any reason to think that "'p' is T" (the left side) is warrantedly assertible is a reason to think that "p" (the right side) is warrantedly assertible.

(2) WA("'p' is T") iff WA("p")

This is not what Wright claims, however: he claims that the deflationist is committed to any reason to think that "'p' is T" is a reason to think that "p" is warrantedly assertible.

(3) "p" is T iff WA("p")

The anti-deflationary argument introduces warranted assertibility on one side of the ES equivalence but not the other. This is not allowed without a further step, one which I assume to be implicit here. In particular, I suggest that Wright takes the deflationist idea of endorsing an assertion to be the idea of ascribing warranted assertibility. So he takes the deflationist's equating the *use* of the truth predicate as a means for endorsing assertions as an equivalence in meaning between endorsing and ascribing warranted assertibility: to say that a sentence or proposition is true is to say that it's warrantedly assertible.

(2a) "p" is T iff WA("'p' is T")

This supplies the missing step from (2) to (3). However, I will argue that Wright, in equating, for the deflationist, the use of "is true" with endorsing assertions, imputes more to the deflationist than the deflationist need accept.

Suppose that Sally asserts, "Today is Wednesday," and Fred, thinking that today is Wednesday, agrees by saying, "That's true." To what is Fred thereby committed? Fred is clearly agreeing that today is Wednesday; hence he is *endorsing* the proposition Sally's assertion expresses. But expressing this as endorsement creates a problematic ambiguity. There are three senses in which we could understand Fred's "endorsement" of Sally's assertion.

(i) Fred is also committing himself to the truth of the proposition asserted (*Today is Wednesday*). Essentially, he's asserting the proposition by using the prosentence "That's true." However, Fred takes no position on the warranted assertibility of the proposition expressed in Sally's assertion.

(ii) Fred is saying that what Sally asserted is warrantedly assertible without thereby saying that *Sally* is warranted in her assertion. Perhaps Fred doesn't think that Sally has sufficiently good evidence for it being Wednesday.

(iii) Fred is saying that what Sally asserted is warrantedly assertible and that *Sally* is warranted in her assertion.

The important question is how best to understand what view of endorsement Wright imputes to the deflationist, and what view Wright requires the deflationist to take in order to run the inflationary argument.

Wright seems to suggest that the deflationist is committed to the use of the truth predicate when applied as a prosentence to assertions in terms of (ii) or (iii). But there are reasons to think that the deflationist doesn't need to take this view, or that, at least, even if some

86

particular deflationists have said as much, deflationism isn't committed to (ii) or (iii).[26] Instead, I suggest, the deflationist's use of the truth predicate is best interpreted as (i). We must separate the *act* of assertion from its (propositional) *content*.[27] Warranted assertibility is a normative evaluation of the *act* of asserting, which is separate from the proposition asserted (the content of the assertion).[28] And so speaking in terms of endorsement without making these distinctions introduces an ambiguity in what it means to endorse an assertion: is one endorsing the act or the content? I think that the deflationist's use refers to the content, and not the act.[29] Although Quine described the truth predicate as a means for semantic ascent (to move from the object language to the metalanguage), the deflationist view that truth is merely a means for endorsing an assertions means that truth is also a means for (merely) semantic *assent*. However, instead of speaking of endorsing an assertion as opposed to endorsing what's asserted (*viz.* the proposition asserted), we should transition away from using endorsement in the latter sense. I thus propose that we reserve "endorse" and its cognates for ascribing warranted assertibility in the sense of (iii) above.

One might object that Wright is only interested in the behaviour of the predicates "is true" and "is warrantedly assertible" and so isn't concerned with explicating warranted assertibility in terms of "is warrantedly assertible for subject S (at time t)." But we should be careful: assertions must be performed by *someone* at some time. All assertions must be appropriately indexed to a subject and a time. I know how many coins are in my pocket, but John doesn't.

---

[26]But insofar as deflationists *have* taken this view, Wright's argument is successful.

[27]A distinction also made by Kvanvig (2009).

[28]As Sosa (2011) would say, the norms of assertion involve *performance* normativity.

[29]Of course, some deflationists *may* have said this, but deflationism as a theory of truth isn't committed to such a view.

Suppose that I have ten coins. John, just guessing, truly asserts, "There are ten coins in your pocket." Knowing that there are ten coins in my pocket, I say, "That's true, but you're not warranted in your assertion." Of course, since I *know* that there are ten coins in my pocket, it's warrantedly assertible *for me* to assert "There are ten coins in my pocket." We're asserting the same proposition, since "your" is indexed to me. But since John lacks adequate evidence for his true assertion, it's not warrantedly assertible for *him.* And, as discussed above, my use of "That's true" doesn't commit me to claiming that John was warranted in his assertion; rather, my use of the prosentence is a way to assent to the truth of the proposition asserted: I make no claims regarding the conditions for warrantedly asserting that there are ten coins in my pocket. Thus, my use of the prosentence is in the sense of (i) rather than (ii) or (iii).

Now, instead of John simply guessing based on no evidence whatsoever, suppose that John hears the characteristic sound of coins in my pocket, and infers that there are multiple coins. So he has *some* reason to think that "There are ten coins in your pocket" is true. But his evidence is so weak (I could just have pieces of metal that aren't coins) that his reasons aren't *good* reasons to think that his assertion is warranted. But since I have adequate evidence, the proposition "There are ten coins in my pocket" is warrantedly assertible, but only *for me* and not John. This is why it's so easy to effect an extensional divergence between "is true" and "is warrantedly assertible," particularly with cases of lucky guesses.

Here's the key to my position: the deflationist's use of the truth predicate to describe assertions is therefore not as a means of endorsing assertions: it's a means of assenting to what's asserted (the propositional content) rather than how or that the proposition was

88

asserted (the act). So we should cease describing the deflationist's use of the truth predicate as a means for endorsing assertions. Ascribing the truth predicate to an assertion is simply to agree with the propositional content of the assertion. Fred says of Sally's assertion that it's true that it's Wednesday: he hasn't yet said anything about whether she made a warranted assertion, or even whether the proposition is warrantedly assertible at all. The distinction is non-trivial, inasmuch as it's possible and sometimes important to assent to the content of an assertion without endorsing the act itself. An assertion might express something true, for example, and yet be imprudent or impolite and so not be warrantedly assertible by anyone in a context: "That's true," one might say, "but you really shouldn't have said that."[30]

By claiming that the truth predicate is merely a means for endorsing assertions, I think that the deflationist is merely committed to "is true" being a means for semantic assent: assenting to the content of the assertion, not the act. So although the deflationist is committed to any reason for thinking "'p' is true" is warrantedly assertible will be a (defeasible) reason for thinking that "p" is warrantedly assertible (since the ES is a claim of semantic equivalence), the deflationist is not committed to "any reason to think that a sentence is T may be transferred, across the biconditional, into reason to make or allow the assertoric move" to assert that p.[31] However, Wright's inflationary argument crucially depends on the deflationist being committed to the truth predicate having the role of assenting to assertions

---

[30]Even if we restrict ourselves to epistemic warrants, the same distinction arises. If Fred knew that Sally was just guessing, then he may assent to what she asserts (because it is Wednesday after all) while still thinking (and saying, if it comes to that) that she made a poor assertion (since she's guessing).

[31]Wright (1992, p. 18). It's important that reason for one can only provide a *defeasible* reason for asserting the other because although two assertions may be equivalent in meaning, pragmatic features inherent in asserting may make one warranted but the other not. "Everything he says is true," for example, is equivalent in meaning to my repeating everything 'he' says, but the pragmatics of that very long speech act make it absurd.

in terms of (ii) or (iii). Without it, Wright's claim that the deflationist is committed to "is true" and "is warrantedly assertible" being normatively coincident is false. And without the claim that "is true" and "is warrantedly assertible" are normatively coincident, Wright's inflationary argument doesn't get off the ground.

And so we may reasonably wonder whether Wright's inflationary argument successfully undercuts deflationism. Again, it's critical to Wright's inflationary argument is that the only factor that explains the normative coincidence but extensional divergence between truth and warranted assertion is some property of truth. Since it's obvious, by now, that "is true" and "is warrantedly assertible" are normatively coincident but extensionally divergent, any adequate theory of truth must account for this. Wright's argument is that deflationism cannot. However, I argue, deflationism is well equipped to do this: it need not be a property of truth that explains the normative coincidence and extensional divergence, but a defining property of warranted assertibility. I've already argued here that assertion plausibly has the goal of truth, and that norms derive their content by what it means to be a rational means of properly aiming at the goals of the practice. That is, warranted assertibility just is "properly aiming at truth." However, it does not yet follow from assuming that truth is the constitutive aim or goal of assertion, that truth is thereby the constitutive norm of the practice. It remains an open question whether one can properly aim at truth while still falling short: that is, whether one may warrantedly assert something false.[32]

As it happens, I think that there are cases where one may properly aim at asserting

---

[32]Since Wright defends a version of the truth norm (in terms of *superassertibility*), this is something that Wright would deny.

truthfully (*i.e.*, warrantedly asserting) while asserting falsely. For this discussion I will borrow from the rich literature of epistemic justification.[33] Epistemic externalists – specifically, reliabilists – take the position that a belief p is justified iff the belief was formed through a reliable (cognitive) belief forming process.[34] What is important for my purposes is how "reliable" is understood; namely, that a process is reliable iff it produces a sufficiently high proportion of true beliefs to false beliefs. What value suffices for "reliable" is not my present concern. What is important is that a fundamental property of epistemic justification is that a belief is justified iff it properly "aims" at being true (*viz.* was formed by a reliable belief forming process).[35] I suggest that the picture is similar for a fundamental property of an assertion being warranted is that it also "properly aims at truth." Aiming at warranted assertion really is aiming at truth. However, this is a defining property of warrant rather than of truth. The fact that warranted assertion is connected to truth is explained by that assertoric warrant means properly aiming at truth; that is, it isn't in virtue of some property of truth that norms of warranted assertion necessarily involve an important role for truth. And, moreover, it doesn't follow from truth being the constitutive aim of assertion that truth is a constitutive norm of the practice. One would need further argument for this claim.

Through this examination of Wright's influential argument against deflationism, we are

---

[33]For example see Alston (1988, 2005), Goldman (1976, 1986), and Plantinga (1993b,a). Since there is a very large body of support for the position that norms of belief and norms of assertion are very closely linked, I find it largely uncontroversial to discuss epistemic justification as an analogue to warrant in warranted assertion. I am referring to what is sometimes called the "belief assertion parallel." See Dummett (1981), Williamson (2000), Adler (2002), and Douven (2006).

[34]Goldman (1976).

[35]In fact, using a particular theory of justification may not be required for my argument. All that is required is that there can be justified false beliefs and unjustified true beliefs. Most theories of justification (*e.g.*, internalism, reliabilism, virtue epistemology, *etc.*) allow for these.

now in a better position to explain the observation that truth and warranted assertibility are normatively coincident but extensionally divergent. Let us begin with the latter. It is clear that there are true beliefs (or propositions) which an agent is not necessarily warranted in asserting. The easy case is lucky guesses. The truth of a lucky guess is not sufficient to provide warrant for assertion. Thus, the extensions diverge because there are some instances of true propositions which are not warrantedly assertible. But what explains this divergence in extension between truth and warranted assertibility is a lack of warrant and not some property of truth. With respect to the observation of the normative coincidence of truth and warranted assertibility I believe that we can find the explanation in the analogous case of truth's role in epistemic justification. That is, a belief is justified iff it "properly aims" at truth insofar as it is the product of a reliable belief forming process. Thus, necessarily, truth plays an important role in norms of epistemic belief (justification) but this is not a property of truth: it is a defining property of justification. Analogously, it is an apparent defining property of warranted assertion that the proposition asserted properly aims at truth (perhaps even that it must be true if truth is a constitutive norm of assertion).

The upshot of this is that what explains the normative coincidence but extensional divergence of truth and warranted assertion is a property of *warrant* and not truth. Thus, I have argued that it is not a property of truth that explains this difference. Since Wright's inflationary argument critically depends on the requirement that it is some property of truth (and not warrant) that explains this difference, I have argued that Wright's argument fails. Truth plays an important role in warranted assertion but what explains the fact that truth

and warranted assertion are normatively coincident but extensionally divergent is a defining property of warrant rather than truth.

Some may claim that an adequate theory of truth must explain the many and various roles that truth plays in our practices. The claim is that deflationism, if it's to be an adequate account of truth, should explain why truth has the normative role in assertion that it does. Since the ES says nothing about warranted assertibility, deflationism can't seem to explain truth's role in norms of assertion. Therefore, the argument goes, deflationism cannot be an adequate theory of truth. However, this line of argument is misguided. We use objects for many different purposes, some of which the product has been designed to fill. A bicycle is used, generally, as a means for transportation. It can also be used as a rather large and expensive paper weight. However, that a bike can be used as a paper weight need not be explained by the design features of the bike.[36] Similarly, a theory of truth doesn't need to explain why and how we choose to use a concept such as truth for our practices, the theory need only explain that it can be used that way. Why truth has a normative role in assertion isn't to be explained by a theory of truth, but rather by our theories of assertion. Can deflationism fill truth's role in the norms of assertion? There seems to be strong *prima facie* support for thinking that deflationary truth is adequate for its role. When we think that "'p' is true" is assertible, we take this as (defeasible) reason to think that "p" is assertible. We do use truth as a means for semantic ascent, assent, and endorsement. And the norms of assertion don't seem to require anything more substantial from the truth predicate than "'p'

---

[36]Obviously, that a bike can serve as a paper weight can be explained by the *physical* features of the bike, but the designer of a bike doesn't need to invoke "bike" concepts to explain this possible use.

is T iff p." Consequently, I think that Wright's inflationary argument against deflationism fails.

This is important in two respects. First, there's a question of whether a deflationary theory of truth is a plausible candidate for an adequate theory of truth. It has been suggested, most forcefully by Wright, that deflationary truth can't handle truth's role in norms of assertion. However, as I have argued, the argument fails. So this attack against deflationism fails. It's still an open question, though, whether deflationism is an adequate theory of truth, and I take no position on the matter. However, the terrain is tricky. There are difficult questions of interpretation involved in my criticism of Wright, some of which may be controversial. But what *really* matters to my project is making clear what it means to endorse an assertion, and the attendant distinction between assenting to the *act* and assenting to the *content* of an assertion. I've argued that, when we say of a speaker's assertion, "That's true," we should normally be interpreted as assenting to the content of a speaker's assertion, and not as endorsing the assertoric act itself. To utter the prosentence is not to express a view on whether the speaker's assertion was warranted, under whatever standards of evaluation should be brought to bear on the speaker's performing that act in context. It is to express a view on whether the *endorser* is, in the immediate context of endorsement, warranted in asserting that content.

Second, Wright is an influential proponent of a truth norm of assertion. Part of his purported motivation for a truth norm is that we use the truth predicate as a means for endorsing assertions. I have suggested that we should separate endorsing an assertion (as

ascribing warranted assertibility) from assenting to the propositional content of an assertion. I think our use of the truth predicate is best interpreted as the latter, rather than the former. Moreover, I've argued that we can understand a central role for truth in assertion, perhaps as the constitutive goal of the practice, without thereby taking truth as the norm of assertion.

4.5   Weiner and the Truth Norm of Assertion

Wright's focus on truth as a norm takes the form of epistemically constrained superassertibility, where a "statement is superassertible...if and only if it is, or can be, warranted and some warrant for it would survive arbitrarily close scrutiny of its pedigree and arbitrarily extensive increments to or other forms of improvement of our information."[37] But the most direct appeal to truth in the context of assertoric norms is simply to claim that truth is the norm of assertion without moreover offering a specific theory of truth. One of the most prominent contemporary theorists in this mould is Matthew Weiner, who defends a straightforward truth norm. As noted previously,

TNA: One should assert that p only if p (is true).

There's a tremendous intuitive connection between warranted assertibility and truth. Williamson notes that it's common to view someone who has asserted falsely as having somehow broken a "rule" of the game of assertion. This is apparent in how we view lying as "cheating." Partly resulting from this intuition, both TNA and knowledge norm (KNA)

---

[37]Wright (1992, p. 48).

advocates take truth of the proposition asserted as a necessary condition for warranted assertibility. As noted previously,

KNA: One should assert that p only if one knows that p.

We can call norms that have truth as a necessary condition *factive* norms. Weiner argues that TNA can adequately address many of the motivations supporting KNA and that TNA can properly account for the propriety of assertions that KNA incorrectly deems unwarranted. He argues that the truth norm, combined with other more general conversational norms (such as Grice's maxims) provides an adequate norm of assertion.

Weiner argues for TNA through some illustrative examples intended to demonstrate that one may warrantedly assert something one doesn't know, provided that one has good reasons to think that it's true, and that it is true. The central case involves a naval captain predicting an upcoming attack.

Captain Jack Aubrey has had long experience of naval combat against the French Navy. He and young Lieutenant Pullings have been watching French ships maneuver off Mauritius all day. At 2 p.m., Aubrey says to Pullings, ["The French will wait until night to attack."]

Weiner suggests that Aubrey's assertion is warranted, provided that the French actually attack at night, but unwarranted if the French do not attack at night. It's taken as *prima facie* plausible that Aubrey does not know, at the time of assertion, that the French will attack at night. However, had Aubrey received a report from a spy indicating that the French were planning to attack, perhaps Aubrey's assertion would rise to the epistemic level

96

of knowledge.[38] Weiner considers it evidence for his view that if the French do attack at night, then Pullings would not be entitled to feel resentment to Aubrey's assertion, but Pullings would be entitled to resent the assertion if the French do not attack.

Insofar as Aubrey's assertion is warranted, KNA cannot account for its propriety; it's stipulated that the assertion does not rise to the epistemic status of knowledge. But if truth is really the norm of assertion, then how does it explain the impropriety of asserting lucky guesses? Lucky guesses are instances where one asserts a true proposition without any (or with very poor) evidence for its truth. If Aubrey had asserted that it's currently raining in Lisbon, without any evidence concerning the weather in Lisbon, then even if it happens to be true, he made an unwarranted assertion. If truth is the norm of assertion, then why isn't this assertion warranted?

Weiner explains the problem of lucky guesses by invoking a distinction made popular by Keith DeRose between primary and secondary propriety.[39] An assertion is primarily proper if it satisfies the norm of assertion. An assertion is secondarily proper if the speaker has good reason to think that the assertion satisfies the norm.[40] Consequently, an assertion can be primarily proper without being secondarily proper, and this is the case for lucky guesses. Since Aubrey would know that he lacks evidence for his assertion about the weather in Lisbon, he can't have good reason to think that his assertion satisfies the truth norm; that is, his assertion lacks secondary propriety. Only assertions which have both primary and secondary

---

[38]I have serious reservations about this, but this is something that is taken up in more detail in Chapter 5.
[39]DeRose (2002).
[40]It's worth noting that this distinction can be made whenever norms of action are in force. The distinction is thus not restricted to norms of assertion.

propriety fully satisfy the norm of assertion. Assertions having the former but not the latter are inappropriate but blameless; those lacking both are blameworthy. So on TNA, a lucky guess is inappropriate but blameless.[41]

On DeRose's view, then, the truth of a proposition is not sufficient for full warranted assertibility, though it is a necessary condition. Speakers must moreover have some evidential basis for the truth of their assertion. What warrants Aubrey's assertion that the French will attack at night is both that it's true (we assume) and that he has good reason to think that it's true: he's an expert in naval combat and he's been observing the French Navy's maneuvers all morning and afternoon. So his assertion, although it's a prediction, is far more than a lucky guess: it's a true assertion for which he has excellent evidence. The difference between this assertion, and one where Aubrey knows, is that there are cases where one's evidence is sufficient for warrantedly asserting something true but fails to rise to the epistemic level of knowledge. Weiner presents Aubrey's prediction as just such a case.[42]

However, it's not clear that Aubrey's assertion is as clearly warranted as Weiner claims. He argues that Pullings will feel entitled to resent the assertion only if it turns out to be false. But why should such a highly contingent feature of the situation – so far beyond Aubrey's control – determine the propriety of the assertion? We can suppose that Aubrey relies on his experience in forming the assertion, "The French will attack at night," and this statement would have been true had it not been for an asteroid striking the Earth and destroying the

---

[41]This is contentious, though. When a speaker has no evidence for the truth of an assertion, then the speaker doesn't deserve any credit (in terms of being not blameworthy) just for getting lucky. Rather, the speaker seems blameworthy for the assertion.

[42]One could object, though, that it's not possible for a speaker to have sufficient evidence for warrantedly asserting a true proposition but insufficient evidence for knowledge.

entire French Navy (including the British, of course) in a massive tsunami. Would Pullings feel entitled to resentment even if what made Aubrey's assertion false were an exceptionally unlikely, and reasonably unforeseen, event of that sort? Intuitions might differ, but insofar as one's intuition suggests that Pullings should not resent an assertion rendered false by a reasonably unforeseeable event, then Weiner's unqualified claim is dubious. In a certain sense, Weiner is enfranchising the resentment of false assertions based on bad luck. *Mutatis mutandis*, this seems to suggest that when Aubrey's assertion is true, it's true partially by luck: luck that his assertion wasn't made false by some event that could not reasonably have been foreseen.[43] And in any case, although one may grant that Weiner can accommodate the unassertibility of simple lucky guesses by appeal to secondary impropriety, there's another pressing problem for the TNA advocate: lottery propositions. While this is a topic treated at length in Chapter 5, a brief discussion of the problem and Weiner's broadly Gricean solution is in order.

It's generally accepted that prior to information on the results of the draw of a fair lottery, it's inappropriate to assert, "Your ticket did not win."[44] If we suppose that there are 1000 tickets in the lottery, then it's *a priori* that this assertion will almost certainly be true (99.9% likely, in fact). And we may further suppose that the assertion is true, but it's intuitively unassertible: how do you know that my ticket did not win? KNA advocates such as Williamson consider such assertions worthy of criticism, particularly with the "How do you know?" challenge. So KNA advocates have taken the lottery problem as evidence

---

[43]Compare Lackey (2008).

[44]Compare, for example, Williamson (2000), Hawthorne (2004), and Hill & Schechter (2007).

against TNA in favour of KNA: what explains the unassertibility must be that a speaker cannot know that a ticket did not win, prior to information on the draw. So if a speaker can assert something true for which she has excellent (though probabilistic) evidence, how can TNA explain the apparent unassertibility of lottery propositions?

Weiner argues that TNA can explain the impropriety of asserting lottery propositions by invoking H.P. Grice's well-known theory of conversational maxims. The two maxims taken as most relevant for Weiner's purposes are the maxims of Quality and Quantity, both of which fall under the general Cooperative Principle.[45]

Cooperative Principle (CP:)

> 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.

Quality:

> Super Maxim: Try to make your contribution one that is true.
>
> Sub Maxims:
>
> 1. Do not say what you believe to be false.
>
> 2. Do not say that for which you lack adequate evidence.

Quantity:

---

[45]Grice (1989), especially Chapter 2.

1. Make your contribution as informative as is required (for the current purposes of the exchange).

2. Do not make your contribution more informative than is required.

To see the maxims in action, suppose that Sarah asserts, "Your ticket did not win," to Alice, based solely on her information about the improbability of any particular ticket's having won. Weiner argues that prior to information on the results of the draw, and based merely on it being highly likely that the assertion is true, Sarah's assertion violates CP. Her assertion conveys the implicature that she has superior information on the results of the draw, which she doesn't in fact possess. If she doesn't have any information besides the inherent improbability of Alice's holding the winning ticket, then she doesn't have any information that Alice doesn't have, so Sarah's assertion, "Your ticket did not win," will foreseeably mislead Alice. If Sarah means to convey only, "Your ticket almost certainly did not win," then asserting, "Your ticket did not win," communicates more information than is required (that the ticket *lost* rather than *very probably lost*) and so violates the second maxim of Quantity. And since Sarah lacks evidence that Alice's ticket actually lost, beyond its merely being very likely, her assertion fails the second sub maxim of Quality.

Weiner argues that assertions such as Sarah's are only plausibly appropriate (*i.e.*, warranted) insofar as at least one of three following conditions is false:[46]

a) Alice is already aware that her ticket is almost certain not to have won,

b) Sarah is not presenting herself as knowing anything that could strengthen Alice's belief

---

[46]Weiner (2005, p. 233).

101

in (a) that her ticket is almost certain not to have won,

c) Alice does not need reminding that her ticket is almost certain not to have won.

If particularly either (a) or (c) is false, then Sarah's assertion has a point and appears warrantedly assertible.[47] Suppose that a prankster has set up a fake lottery draw to make it appear to Alice that she won. Sarah knows of the plot and when Alice views the tape and thinks that she won, Sarah asserts, "Your ticket didn't win; that's a videotape of a fake broadcast." Weiner considers this a warranted assertion: it communicates something important, and is something Alice doesn't already know.

There is considerable appeal to this way of addressing the problem that lottery propositions pose for TNA advocates. Unfortunately, the argument appears to prove too much; such assertions never seem assertible with full propriety if we take the proposal literally. For example, in the last scenario Sarah has asserted "Your ticket did not win" – which she cannot know. All she knows is that Alice is watching a faked broadcast. Sarah doesn't have information on the actual results of the draw so, for all Sarah knows, Alice's ticket did win in the real draw. This would simply be a special case of the original unassertible proposition. Weiner considers this objection but he doesn't make clear whether he considers Sarah's assertion to Alice in the videotaped drawing case to be warranted all things considered. I return to this problem for TNA in Chapter 5 where I discuss the implications of the lottery paradox in greater detail. I will argue that there are some lottery propositions which seem to pass Gricean maxims, but remain intuitively unassertible, and so constitute an undiminished

---

[47]Weiner (2005, p. 235).

challenge to TNA.[48]

There remains a final class of problematic cases for TNA to consider, though. I've already discussed the simple case of lucky guesses and the TNA advocate's response above, but there is a more serious kind of lucky guess scenario, in which the speaker thinks that she has good evidence for a true proposition, but her true assertion is really based on faulty reasoning. Suppose that Bill is a bad gambler who loves to play roulette. He only places red or black bets (that is, betting on whether the result will be red or black, which has slightly less than 50% probability of winning).[49] He has a "system" according to which he watches the last 20 spins and then bets based on which colour was under-represented in the sample. He thus glaringly commits the gambler's fallacy: when events are independent, one bets on the under-represented events expecting things to "even out" in the end. So if the wheel produced 16 red in the last 20 trials, Bill bets black thinking that it's "bound to come up": black is "due." So he says to his friend, "It's going to be black next, so I'm betting black."

Since Bill knows that he should only place bets where the odds are favourable, in order for his assertion to be (as DeRose and Weiner would say[50]) secondarily appropriate, he must at least think that he's satisfying the central norm of assertion. If we suppose that TNA is the norm, then Bill has to think that he's asserting something true and that he has adequate evidence for its truth. He clearly thinks this: he has a "system" after all. Unfortunately, DeRose's specification of secondary propriety, and Weiner's use of it, can't

---

[48]See 5.7 on page 143.

[49]It's a sucker bet, of course, because the bet only pays 1:1 where the odds of winning are less than 1:1, which is required to make the bet have at least a neutral expected value.

[50]Or, as I would say, normatively proper.

distinguish between cogent and faulty reasoning in a speaker's arriving at an attitude towards his own assertion.[51] If we suppose that Bill's bet is successful, then he asserted something true: the next spin did come up black. So his assertion is primarily appropriate. He moreover thinks that his assertion satisfies the norm, and he thinks that he has excellent evidence for the truth of his assertion, so his assertion seems to be secondarily appropriate too. But, of course, Bill's assertion isn't warranted: although he thinks that he has good evidence for the truth of the assertion, he doesn't actually have good evidence. Since the primary/secondary propriety explanation of the impropriety of this lucky guess can't explain the unassertibility of Bill's assertion, it's unclear how the TNA advocate can explain away this case of an intuitively bad assertion.

It's possible, of course, for the primary/secondary propriety distinction to be refined in order to express that an act is secondarily appropriate only if the speaker has *good* reasons to think that her assertion satisfies the norm in the primary sense. However, how shall we specify "good" reasons? Shall they be objective? One of the possible explications of the primary/secondary propriety distinction is that whether one's act is primarily appropriate is an objective fact of the matter: one's assertion either does or does not satisfy the norm in the primary sense. But whether one's assertion is secondarily appropriate is sometimes specified as a subjective fact: whether the agent thinks that she has (good) reason(s) to think that her assertion satisfies the norm. The betting example is meant to show that there are cases where agents can have reasons which appear good by their own lights for

---

[51]One may suppose that DeRose could be an externalist about "having good reason" and so think that one's assertion satisfies the norm of assertion, but he's not clear on this point.

thinking that their assertion satisfies the norm, but are actually poor reasons, and that their assertion therefore seems secondarily appropriate. Bill takes himself to be reasonable in thinking that his assertion is true: he subjectively thinks that he has good reasons for his assertion. But if we are to understand secondary propriety in an objective sense, just as we are to understand primary propriety, then Bill's taking himself to be reasonable will not be sufficient for his assertion being secondarily appropriate. Since his reasons were in fact poor, whether he considers himself to have good reason to think that his assertion satisfies the norm is irrelevant to whether he actually has good reasons to think as such. Unfortunately, none of this is made explicit by DeRose or proponents of the primary/secondary propriety distinction. And so more needs to be said.[52] In Chapter 7 I return to the plausibility of the primary/secondary propriety distinction.

4.6 Conclusion

The principal goal of this chapter was to better understand the relationship between norms of assertion and truth. Some have argued that truth is, in fact, the norm of assertion: assert that p only if p. Under this sort of factive norm, any false assertion will count as violating the norm. Evidence that truth is the norm of assertion is sometimes adduced by noticing that we use the truth predicate to endorse assertions. We say of someone's assertion, "That's true," and this seems to be a means of ascribing warranted assertibility to the assertion. Moreover, we often tend to rebuke or resent false assertions. Williamson has referred to this intuition as

---

[52]Unfortunately, this discussion would take us too far afield, although I consider it an interesting and important question.

being akin to speakers of false assertions somehow having cheated at the game of assertion.

Crispin Wright takes the role of truth in norms of assertion to have important implications for theories of truth. I discussed Wright's influential inflationary argument against a particular theory of truth, deflationism. The argument is that the meaning of applying the truth predicate to assertions leads to a problem for the deflationist. Deflationism is committed to two theses: that saying "That's true" of an asserted sentence amounts to the same semantic content as repeating the assertion; and that there is nothing to the concept of truth and the use of the truth predicate beyond this disquotational function.[53] Wright's argument amounts to claiming that the deflationist is committed to truth being the norm of assertion. And since we notice that there are some true assertions that are not warrantedly assertible, deflationary truth cannot account for this without resorting to a more substantial conception of truth than the deflationist is permitted. Deflationism, then, inflates under pressure.

I have argued that this argument is unsuccessful. The argument seems to trade on a problematic ambiguity in what it means to endorse an assertion. If the deflationist's view of the use of the truth predicate, in terms of endorsing an assertion, amounts to merely assenting to the semantic content without thereby being committed to viewing the assertion as warrantedly asserted, then Wright's argument doesn't work. The argument requires that the deflationist is committed to the use of the truth predicate as ascribing warranted assertibility. But there's nothing in the deflationary concept of truth and the truth predicate that commits the deflationist to this view. The deflationist is *not* committed to truth being a (or

---

[53]Though, as discussed earlier, for a prosentential deflationist there will be pragmatic differences between asserting the prosentence and asserting the sentence.

106

the) norm of assertion. Deflationism is merely committed to the use of the truth predicate when applied to assertions being semantically equivalent to merely repeating the assertion (there may be pragmatic differences, though). To claim that an assertion is true is merely to give semantic assent: one agrees that the proposition asserted is true, and this can be done whether one thinks that the assertion is warranted or unwarranted.

I then discussed the plausibility of truth being the constitutive norm of assertion. Matthew Weiner argues that the central motivations supporting the Knowledge Norm are easily accommodated by TNA combined with general conversational (*viz.*, Gricean) norms. Moreover, he argues, KNA gets some cases wrong, such as Aubrey's prediction, that TNA gets correct. I raised three objections to TNA: two forms of lucky guesses, and lottery propositions. It's plausible that TNA can appeal to the primary/secondary propriety distinction in dealing with simple lucky guesses, but it's not clear how this strategy is available for dealing with more complicated cases where the speaker asserts a lucky guess which satisfies the TNA norm and the speaker seems to have secondary propriety. I discussed Weiner's Gricean explanation of how TNA can accommodate the general unassertibility of lottery propositions such as "Your ticket did not win." However, there may be lottery propositions which satisfy TNA but are intuitively unwarranted, where a Gricean explanation is unavailable. In the next chapter I take up the issue of the assertibility of lottery propositions at length and will return to the question whether TNA can actually well explain the propriety, or impropriety, of such assertions.

Chapter 5

The Lottery Paradox and Knowledge

5.1 Introduction

Statements about lotteries raise parallel problems for epistemologists who want to articulate

conditions for knowledge and those working on norms of assertion.[1] In both cases a central

concern is whether one can know or warrantedly assert sentences like

(A) Ticket n will lose

before the draw is made.[2] A growing consensus has formed that (A) and propositions like it

are neither knowable nor assertible. Some writers on the topic, such as Timothy Williamson,

have argued that this fact is best explained by, and thus provides evidence for, the Knowledge

Norm of Assertion (KNA). Others, such as John Hawthorne, have furthermore argued that

ordinary propositions like *I will not have a heart attack in the next minute* and *My car is*

---

[1]Consider Kyburg (1961), Harman (1973), Vogel (1990), Lewis (1996), Williamson (2000), Hawthorne (2004), Kvanvig (2009), among others

[2]I will in general use 'assert' to mean 'properly assert' or 'warrantably assert' unless otherwise noted.

in the driveway are, properly speaking, lottery propositions.[3] That is, they are relevantly similar to (A) with respect to their knowability and assertibility given probability of less than 1 on one's evidence. An implication of the latter position is that since knowledge is not possible in the traditional lottery case, it is also not possible in the more general cases of everyday life. At least, that is the implication unless a principled line can be drawn between ordinary propositions and traditional lottery propositions.[4]

This chapter is composed of two main projects. In the first part I present the problem for knowledge and assertibility raised by Williamson and Hawthorne's analysis of the lottery paradox. Williamson's presentation of the argument indicates an inference from the unassertibility of lottery statements, to their unknowability, to merely probabilistic evidence being always insufficient for knowledge. As others have noted, Williamson's position, particularly when combined with Hawthorne's, implies something close to an error theory about knowledge and warranted assertion. I will argue that this consequence is both powerfully counterintuitive on its own terms and represents a methodological tension for Williamson inasmuch as his analysis uses *intuitions* about the unknowability and unassertibility of lottery propositions as data.

In the second part I will offer an analysis of the lottery paradox. This analysis, I argue, accommodates the intuitions invoked by Williamson, to the effect that one does not know and may not assert (A) in the lottery case. And, unlike Williamson's account, it does not violate equally strong intuitions about knowledge in the ordinary case. I draw on a family of

---

[3]Hawthorne (2004) gives credit to the initial argument by Vogel (1990).

[4]Of course, both Hawthorne and Williamson are not sceptics: they each have an account detailing how we can come to know much of what we intuitively think we know.

related proposals defined by their sharing the idea that knowledge or assertibility tolerates probabilities less than unity, when the space of possible outcomes includes some that can be properly ignored from consideration. What it is to *properly* ignore some possibilities when considering defeater conditions for knowledge and assertibility is a point on which different theories may disagree; but in broadest outline this family of approaches has the virtue of allowing us to explain an intuitively clear asymmetry between "lottery" propositions and ordinary propositions (which I will call "fallible" propositions). I argue that the reason one cannot know or assert a lottery proposition is that one lacks sufficient epistemic justification for such propositions.

There are a number of implications for norms of assertion. First, my thesis means that the lottery paradox no longer uniquely supports KNA rather than a reasons-based norm. Second, in Chapter 4 I raised the spectre that the lottery paradox may pose a problem for TNA advocates. I will return to that question in more detail. Prior to an adequate analysis of the unassertibility of lottery propositions, such propositions seem to meet Weiner's criteria of well justified true beliefs. But if lottery propositions are paradigmatic cases of unassertible propositions, this is a particular problem for the TNA advocate. However, my analysis of the assertibility of lottery propositions will offer an escape for the TNA advocate. The ultimate upshot is that the pathology of the unassertibility of lottery propositions seems best identified with a lack of sufficient epistemic reasons for the assertion, which is particularly well explained by a reasons-based norm. Therefore, the unassertibility of lottery propositions isn't best explained by appealing to a speaker's lack of knowledge; instead, it's best explained

by appealing to a speaker's lack of adequate epistemic reasons. Consequently, the lottery paradox linguistic data seems to better support a reasons-based norm rather than KNA, contrary to the arguments of KNA advocates.

5.2   The Lottery Paradox

I begin by presenting the lottery paradox. Suppose that there is a fair lottery with n tickets, say 1000, where one and only one ticket will win.[5] The probability of a given ticket losing is 99.9%. Suppose that these grounds are sufficient to rationally believe that one's ticket will lose. Jane holds one of these tickets and on these grounds comes to believe *B1*.

B1: My ticket will lose.

By parity of reasoning, if Jane is justified in believing B1, she will be justified in believing *B2*, for any ticket n.

B2: Ticket n will lose.

If we suppose some reasonable formulation of the closure principle, then Jane will be justified in believing *B3*.

B3: All tickets will lose.

---

[5]In fact, the paradox will result even if it is a biased lottery or even if it's likely that no tickets will win. Williamson (2000, p. 248), DeRose (1996), Douven (2006, p. 462, n18), and Hawthorne (2004, pp. 8-9 and 15). Furthermore, even if it's possible that no ticket will win, Jane will have very good evidence that the proposition *All tickets will lose* is false such that the paradox still results.

However, if this is a fair lottery and one ticket must win, then the probability of B3 is 0 and it's therefore not rational for Jane to believe it. Thus, Jane is both justified and not justified in believing B3, producing a paradox.

This has been a problem for epistemology ever since its first formal presentation by Kyburg.[6] However, I do not wish to enter into the discussions of various attempts to resolve the paradox.[7] Instead, I raise it as a foil for the role of the lottery paradox in the literature on norms of assertion. There is an emerging consensus that prior to information on the results of a draw one is not warranted in asserting *A1*.

> A1: Your ticket will lose.

While there is no consensus on how to best explain this intuition, a prominent attempt belongs to Timothy Williamson. I turn now to his analysis.

5.3   Williamson and the Lottery Paradox

Williamson begins his analysis from the observation that there is a widely shared intuition in our linguistic practices that assertions such as A1 are inappropriate; that is, not warrantedly

---

[6]Kyburg (1961).

[7]One might think that any reasonable formulation of the closure principle that gives us B3 is therefore *un*reasonable. The difficult question of whether *any* closure principle is truly reasonable, all things considered, is not one that I will consider in much detail. However, in my opinion, a promising response is to deny the principle of closure on which an agent justified in holding both conjuncts is thereby justified in holding their conjunction. Kyburg (1961) who originally formulated the paradox suggested just such a strategy for resolution. Hill & Schechter (2007) recognize that aforementioned closure principle will "aggregate risk,'" and that when one conjoins two probabilistic statements, the conjunction will necessarily have a lower probability than both conjuncts. Thus two justified beliefs which individually meet the standard of justification may fail to meet that standard when combined. This would dissolve the paradox. Williamson (2000, p. 248) appears to concur.

assertible. Call this intuition *INT1*.[8]

> INT1: One may not assert that a ticket will lose in a lottery made prior to information on the result of the draw.

One's evidence in the case of lottery propositions such as B1/A1 is that it is highly probable that they are true. Suppose that Ben were to make the flat-out assertion A1 to Jane. If Jane were to find out that Ben's evidence for his assertion was merely that it is highly probably true, then even if it happens to be true that Jane's ticket will lose, she is entitled to regard Ben as having spoken improperly. Her taking the view (and perhaps her accompanying resentment!) reflects an intuition that asserting A1 based merely on its being very probably true is tantamount to cheating at the practice of assertion: Ben exceeds his evidential authority when he asserts A1. The nearest thing to A1 that Ben can legitimately assert is the explicitly qualified *A2*.

> A2: It is highly likely that your ticket will lose.

Linguistic evidence for these claims is that the response "But you don't know that!" is a decisive criticism of A1 but not of A2.[9]

A prominent feature of the lottery paradox is that the particular probability of the truth of A1 is not at issue. The intuitions about assertibility canvassed above seem to survive on any lottery, no matter how large the number of tickets. So we can make the prior probability

---

[8]Williamson is not alone in suggesting that this is a common intuition. See, for example, Hawthorne (2004), Douven (2006), and Kvanvig (2009). I share the intuition that one does not *know* in the lottery case (INT2), but not that one may never warrantably assert (INT1). However, for the purposes of this chapter, I assume that such assertions are never warranted.

[9]Williamson (2000, pp. 247-8). He argues that "probabilistic evidence warrants only an assertion that something is probable'" (p. 248).

that a given ticket will lose arbitrarily close to certainty without undermining the intuition that an agent may not assert A1 (*viz.*, INT1) and cannot know B1. Call this intuition *INT2*.

> INT2: One does not know that a given ticket will lose, in the absence of information on the result of the draw.

Williamson concludes that merely highly probable evidence, even when the proposition is true, does not suffice for knowledge or warranted assertibility.[10] Call this the *Insufficiency Principle* (*INS*).

> INS: That p is true and that p has (arbitrarily) high probability <1 on one's evidence are insufficient for one knowing that p.

This principle generalizes for Williamson beyond the case of lotteries. "What matters in the original lottery case is not the subject matter of the assertion but the probabilistic basis on which it was made," he writes.[11] When one's evidence for the truth of a (true) proposition is merely probabilistic (even if it is arbitrarily close to certain), such evidence is always insufficient for assertibility and knowledge.

Williamson's analysis of the lottery paradox differs importantly from the epistemological analysis presented in Section 5.2. The epistemological analysis requires some principle of closure; however, Williamson denies that justification is closed (at least in these circumstances). He offers an analogy.[12] Suppose that a large pile of chocolates sits in front of me. It's a

---

[10]Below I will briefly return to discuss Williamson's view of evidence and evidential probability on which knowledge that p is probability 1 on one's evidence. I don't mean to here suggest that Williamson arrives at his view of evidential probability *because* of his views of the lottery paradox, though it's possible that this is the case.

[11]Williamson (2000, p. 250).

[12]Williamson (2000, p. 248).

large pile, but one that I have a strong desire to completely consume. However, I have been informed that one chocolate has been poisoned. It will make me ill but I will recover. If we suppose that the pile is very large (and I am very hungry), "I can quite reasonably eat just one, since it is almost certain not to be contaminated, even though, for each chocolate, I have a similar reason for eating it, and if I eat all the chocolates, I shall eat the contaminated one, and my sickness will be overdetermined." However, "no plausible principle of universalizability implies that, in the circumstances, any reason for taking one chocolate is a reason for taking them all; the most to be implied is that, in the circumstances, any reason for taking one chocolate is a reason for taking any other chocolate instead."[13] Williamson thus denies that the justification to take a chocolate, although such justification is present for taking any single chocolate, is closed such that if I am justified in taking one chocolate, then I am justified in taking two chocolates . . . then I am justified in taking all chocolates, which is absurd.[14] Yet the intuition remains that even though I may be justified in taking a chocolate, I cannot be said to *know* that a given chocolate is uncontaminated. The chocolate case is analogous to the lottery case since, *mutatis mutandis*, INT1 and INT2 apply here also. I may not assert that a given chocolate is uncontaminated (before I eat it) and I do not know that a chocolate is uncontaminated (before I eat it).

---

[13]Williamson (2000, p. 248). Set aside the empirical disanalogies between lottery cases and the diminishing appeal of eating chocolate as one eats more of it. Williamson's analysis of the chocolate-eating case parallels the epistemological analysis of the lottery paradox, on which the justification for believing that any one ticket will lose applies equally to all other tickets, and closure seems to license the move from believing of each ticket that it will lose to the obviously false belief that all tickets will lose. Believing that all tickets will lose is analogous to its being overdetermined that I will become sick from eating the contaminated chocolate.

[14]Others have suggested a similar tactic to the lottery case in both the epistemological literature and the assertion literature – see note 5.

How, then, do these various intuitions and principles fit together? Williamson begins with the widely shared intuition that assertions such as A1 are inappropriate (INT1). There is linguistic data supporting this. An intuitively acceptable response to such assertions is to object that the speaker does not *know* that the ticket will lose (INT2). The intuition is that, as far as the speaker knows, the ticket might win; that is, since the assertion is made before information on the outcome of the draw, the speaker does not know that the ticket will lose. Williamson is a central figure who argues that the constitutive norm of assertion is the Knowledge Norm of Assertion (KNA).

KNA: One may assert that p only if one knows that p.

Since he thinks that knowledge is the norm of assertion, Williamson thinks that the best explanation of why assertions such as A1 are not warrantedly assertible is that the speaker lacks knowledge. Williamson observes that the speaker's evidence for A1 is merely that it is (arbitrarily) highly probably true; moreover, even when it is true that the ticket will lose, the speaker lacks knowledge. So the best explanation for why the speaker lacks knowledge is that merely highly probable evidence, even when true, is *never* sufficient for knowledge.[15] So, Williamson uses INT2 to explain INT1 and then produces the INS principle.

Of course, Williamson's commitment to merely probabilistic evidence being insufficient for knowledge may be independent from his treatment of the lottery paradox. For Williamson, knowledge is evidence. Moreover, all and only knowledge is evidence. Thus, E(vidence) = K(nowledge). From this, he argues, knowledge that p is probability 1 on one's *total evidence*.

---

[15]Williamson (2000, p. 250).

If E = K, then part of my evidence for my knowledge that p is my knowledge that p. So even if p(e) = 0.1, p(e|e) = 1. And if I know that p, then p(p|e) = 1, since part of my total evidence is my knowledge that p (from E=K). That is, when I know that p, that knowledge is added to my set of evidence for p. And so we can see how it follows that merely probabilistic evidence is insufficient for knowledge: conditionalizing on merely probabilistic evidence doesn't produce the result that p(p|e) = 1.

This argument raises some concerns, though. One concern is the conclusion that merely highly probable evidence, even when the proposition is true, is never sufficient for knowledge. I will argue that while this is true in the case of lotteries, it's not true for many ordinary propositions such as "I'm holding my coffee cup." Another concern is that this may commit Williamson to a strongly counter-intuitive account of knowledge and warranted assertibility claims and ascriptions. The problem comes from his argument using an intuition as a datum. If using an intuition leads to a strongly counter-intuitive and revisionist account, then this casts doubt on the evidentiary status of the initial intuition which was used as a foundation for his argument. However, I only raise these concerns now in order to set them aside until Section 5.5.

5.4 Hawthorne and Widening the Sense of Lottery Propositions

Following Vogel, John Hawthorne argues that we should consider a wide class of propositions to be lottery propositions.[16] Lottery propositions share two characteristics: they are (arbitrarily) highly probable on an agent's evidence and there is an intuitive reluctance to say that one can know or warrantedly assert them. Examples abound. For example, A3 is a lottery proposition.

A3: I will not have a heart attack in the next 5 minutes.

Suppose that I am a healthy adult and I have just finished undertaking a rigorous physical exam by the best team of doctors in the world. They assure me that my health is excellent. In these epistemic circumstances, we may fairly describe the probability of my suffering a heart attack in the next five minutes as exceptionally low. Yet Hawthorne argues that we can reconstruct the case as my holding a lottery ticket where, oddly enough, "winning" is having the heart attack. Since it is very unlikely that I will "win," the lottery has many tickets. Let us suppose that there is a one in a million chance of my suffering the heart attack. But even if the lottery has one million tickets, Hawthorne argues that the intuitions INT1/INT2 remain even if it is true that I will not suffer the heart attack. For, INT1/INT2, by Williamson's argument, arise even if the truth of the proposition is arbitrarily close to certain and even when it is true. Thus, Hawthorne argues that if a proposition can be properly

---

[16]Hawthorne (2004); referring to Vogel (1990).

reconstructed to be the outcome of a lottery, then one can neither know nor warrantedly assert the proposition.[17]

The key to Hawthorne's proposal is parity reasoning. In the original lottery case, Jane doesn't have any more reason for thinking that her ticket will lose than she has for thinking that any other ticket will lose. So at the root of lottery propositions is that they can be conceptualized as being divided into a "possibility space of a set of subcases, each of which, from the point of view of the subject, is overwhelmingly likely to not obtain, but which are such that the subject's grounds for thinking that any one of the subcases does not obtain is not appreciably different than his grounds for thinking that any other subcase does not obtain."[18] He goes on to write,

> In general, what is often at the root of the relevant lottery intuition is a division
> of epistemic space into a set of subcases with respect to one's epistemic position
> seems roughly similar. Once such a division is effected, a parity of reasoning
> argument can kick in against the suggestion that one knows that a particular
> subcase does not obtain, namely: If one can know that *that* subcase does not
> obtain, one can know of each subcase that it does not obtain. But it is absurd to
> suppose that one can know of each subcase that it does not obtain.[19]

Hawthorne thus proposes the following principle of *Parity Reasoning*.[20]

---

[17]Of course, Hawthorne attempts to argue for a principled distinction between cases that can, in principle, be reconstructed as lotteries but where knowledge is nevertheless possible. It's a "salience" version of a relevant alternatives theory (RAT). I discuss such attempts beginning in Section 5.5.

[18]Hawthorne (2004, pp. 14 - 15).

[19]Hawthorne (2004, p. 15).

[20][p. 16]hawthorne04.

Parity Reasoning. One conceptualizes the proposition that p as the proposition that one particular member of a set of subcases $(p_1, \ldots, p_n)$ will (or does) not obtain, where one has no appreciably stronger reason for thinking that any given member of the set will not obtain than one has for thinking that any other particular member will not obtain. Insofar as one reckons it absurd to suppose that one is able to know of each $(p_1, \ldots, p_n)$ that it will not obtain, one then reckons oneself unable to know that p.

And so we can see how a proposition such as *I will not have a heart attack in the next 5 minutes* can be reconstructed as a lottery proposition such that Parity Reasoning can kick in. Call this proposition H. Divide up the possibility space for my having a heart attack in the next 5 minutes. Suppose that the probability of my having a heart attack is vanishingly small, but nonzero. Relative to my epistemic position, my reasons for thinking that, in this case, H is true are not appreciably different than in cases where H is false. Therefore, by Parity Reasoning, if I can't know of *each* and every subcase that H is true, I cannot know in this subcase that H is true. The upshot of Hawthorne's argument is that a proposition can be reconstructed as a lottery proposition when the proposition can be conceptualized where the possibility space can be divided such that Parity Reasoning applies.

Now, it's natural to think that lottery propositions only cover instances of prediction, since there's an inherent uncertainty about the future. However, another of Hawthorne's examples suggests that many propositions about the present are also lottery propositions. Consider A4.

A4: My car is parked in the driveway.

Suppose that I have just looked to my driveway and clearly saw my car through my front window. My car really is in the driveway. Normally we take this to be sufficient to count as something that I know. But suppose that before I make the assertion A4 I turn around so that I no longer see my car: do I still know? It's at least *possible* for my assertion to be false: given my evidence about the location of my car just a moment ago, it's possible for it to have been stolen. Of course, given my background understanding of how long it takes for someone to steal my car, I consider the possibility extremely remote (even nearly impossible). And we can reconceptualize the proposition into a possibility space where the proposition is true in the vast majority of subcases, but my reasons for thinking that A4 is true in this subcase is not appreciably different from subcases where A4 is false. So, pace Hawthorne, this is still a lottery proposition and INT1/INT2 should remain even if it's the case that my car really is in my driveway (*viz.*, the proposition in A4 is true). Therefore, unless a distinction can be made to save the knowability of propositions like A4, Hawthorne's argument implies that we cannot know many of the propositions that we ordinarily take ourselves to know.[21]

While the car example demonstrates that lottery propositions range over claims about the present, the pervasiveness of lottery propositions extends even into statements about the past. Consider A5.

A5: The Magna Carta was signed in 1215.[22]

---

[21]Of course, Hawthorne later proposes a salience version of a relevant alternative theory in order to save our ordinary claims to knowledge. See 5.5.

[22]One might think that there's a difference between this case and the previous cases. The former all seem to be both epistemically possible and physically possible, whereas the latter is merely epistemically

We may have very strong evidence for the truth of the proposition in A5 such as explanatorily coherent and mutually confirming separate accounts that indicate the year as 1215. Still the possibility remains that all of these accounts are mistaken. So, Hawthorne would suggest, we could consider ourselves as holding a ticket to the *The-Magna-Carta-was-not-signed-in-1215* lottery. We expect to lose since it is unlikely that the Magna Carta was not signed in 1215 but it is possible that it was not. And we can reconceptualize the proposition such that Parity Reasoning would apply. Since it's possible that we could "win" this lottery, and our reasons for thinking that A4 obtains in this subcase are not appreciably different from cases where A4 is false, Hawthorne would argue, INT1 and INT2 suggest that we do not know and may not assert A4.

Lottery propositions, therefore, range over instances of predictions, retrodictions, statements about the present and perhaps others. They pervade the sorts of propositions which we intuitively take ourselves to know and warrantedly assert. I may have very good evidence that I'm holding my favourite coffee cup at the moment, but if I close my eyes and assert "I am holding my coffee cup" I must admit that it's possible that someone has quickly and surreptitiously switched my cup for something that has the identical weight and feel of my cup. Call this assertion A6.

A6: I am holding my coffee cup.

Since, considering the possibility space, my reasons for thinking A6 true are not appreciably different in subcases where A6 is false, A6 is a lottery proposition.

---

possible (if we assume that the Magna Carta *was* signed in 1215). For epistemologists such as Williamson and Hawthorne, these cases only need to be epistemically possible to produce the relevant problem. See Hawthorne (2004, pp. 111-112, and elsewhere) where he discusses the Epistemic Possibility Constraint.

Hawthorne's analysis of the generality of the lottery paradox is an extension of Williamson's analysis. Hawthorne argues that propositions that can be reconstructed as lotteries, and where Parity Reasoning applies, are lottery propositions. The important implication is that when we can properly perform this reconstruction, one's evidence is "merely probabilistic" in Williamson's sense. Thus, insofar as it is not the subject matter *qua* lotteries that makes such propositions unknowable (and not warrantedly assertible) but rather that one's evidence is merely probabilistic, Williamson's position appears to lead to Hawthorne's.

A striking implication of these arguments is that we tend to be radically mistaken about what is knowable and assertible. This would suggest something close to an error theory or skepticism about knowledge and warranted assertibility claims and ascriptions.[23] We are just generally mistaken when we think that we know or may assert something. I'm confident in my intuition that I know and may warrantedly assert that I am holding my coffee cup, even immediately after I close my eyes. I have good reasons to think so, despite recognizing the possibility that I am mistaken and therefore recognizing that I could reconstruct the case as a lottery where Parity Reasoning would apply. The arguments of Hawthorne (and Williamson) suggest that when my intuitions conflict with the results of this analysis of lottery propositions, then so much the worse for my intuitions. This strong counter-intuitiveness is a problem shared by Williamson and Hawthorne.

---

[23]Douven (2006) and Dodd (2007), among others, have also noticed this potential consequence of Williamson's position.

5.5 An Alternative Analysis of the Lottery Paradox

Let's assume for the sake of argument that in the lottery case (B1/A1) one does not know despite having (arbitrarily) highly probable evidence for a true proposition. It is reasonable to conclude that having a highly probable true belief is not always sufficient evidence for knowledge. However, it is an entirely different matter to think that having a highly probable true belief is never sufficient for knowledge. Since nearly any proposition can be reconceptualized as a lottery proposition (or one that entails a lottery proposition), the suggestion that highly probable evidence is never sufficient for knowledge is what opens the door to the error-theoretic implications of Williamson's and Hawthorne's arguments. Thus, even our ordinary knowledge claims seem based on merely highly probable evidence. It remains to be seen whether a principled line can be drawn between "ordinary" and "lottery" cases in order to save the intuition that we often do know in the ordinary case (an intuition both Williamson and Hawthorne want to retain). One strategy is to establish that in some class of cases merely highly probable evidence for true propositions is sufficient for knowledge provided that the agent has the right sort of merely probabilistic evidence. I turn now to provide such a proposal.

First, I suggest making a distinction between what have come to be called "lottery propositions" in the literature and what I will call "fallible propositions." I suggest that we reserve the name "lottery propositions" for propositions about genuine lotteries. In this sense A1 is a lottery proposition but A6 is not. Let us call "fallible propositions" those propositions that Hawthorne argues may be reconstructed as a lottery but are not actually about lotter-

ies. In this sense A3-6 are fallible propositions. This distinction will facilitate establishing a principled difference between the two kinds of propositions pace Hawthorne.

I argue that a promising strategy for providing a principled distinction between lottery and fallible propositions can be found in a family of related proposals sharing the idea that knowledge or assertibility tolerates probabilities less than unity, when the space of possible outcomes includes some that can be properly ignored; that is, properly excluded from consideration. What it is to *properly* ignore some possibilities when considering defeater conditions for knowledge and assertibility is a point on which different theories may disagree; but in broadest outline this family of approaches has the virtue of allowing us to explain an intuitively clear asymmetry between lottery and fallible propositions. I will refer them as "relevant alternative" theories (RATs).[24]

An established example of a RAT is David Lewis' version of contextualism. That it is a contextualist account is not important; there are also non-contextualist RATs, after all. Lewis appeals to possible worlds in explicating the concept of alternative possibilities, and proposes that "*S knows* that P iff S's evidence eliminates every possibility in which not-P – Psst! – except for those possibilities that we are properly ignoring."[25] The concepts of *eliminating* and *ignoring* alternative possibilities, and not the specifically Lewisian view of possible worlds, are what make this a clear case of a RAT.[26] More generally: in order for

---

[24]I use this term with some reluctance because there is already a distinction in the literature between specific theories known as, for example, *the* Relevant Alternative Theory (and the "New" Relevant Alternatives Theory), and Lewis' contextualist account as well as non-contextualist accounts. I think that the family resemblance is so strong that it is reasonable to refer to the whole group as relevant alternative theories since they all focus on the feature of dividing the possibility space into relevant and irrelevant alternatives. See Vogel (1999), for example.

[25]Lewis (1996, p. 378).

[26]For a good discussion, see Hendricks (2006). I do not wish to wade into the debate over what fixes the

Subject S to know that p, S must eliminate all of the relevant alternative possibilities where $\neg p$ while ignoring irrelevant alternative possibilities where $\neg p$. [27] Call this the RAT principle.

> RAT: In order for Subject S to know that p, S must eliminate all of the relevant
> alternative possibilities where $\neg p$ while ignoring irrelevant alternative possibilities
> where $\neg p$.

This is a principle common to all relevant alternative accounts. In order to eliminate a relevant possibility, $\neg p$, an agent will need evidence for $\neg\,\neg p$. For example, suppose that I see a yellowish bird in a garden and a friend asks me whether there's a canary or a raven in the garden.[28] I respond "A canary" because, given the only two relevant possibilities for the species of the bird, and since it is yellow and not black (the colour of a raven), I have eliminated the possible world where $\neg p$ *viz.* the raven-in-the-garden world.[29] However, if the question were "Is there a canary or a goldfinch in the garden?" then I would not have sufficient evidence to assert that there is a canary since I do not have evidence that it is not a goldfinch. In this case I do not have sufficient evidence to eliminate the alternative possibility where $\neg p$ *viz.* the goldfinch-in-the-garden world.

Providing a precise definition of what constitutes relevant alternatives requiring elimination, or irrelevant alternatives which one is permitted to ignore, is not an easy matter. This is largely where disagreement between various RATs is found. Lewis himself fails to

---

context; for example, whether it's speaker- or ascriber-dependent. This is a notoriously difficult topic which I believe does not have a significant bearing on my use of the contextualist account of knowledge.

[27]Of course, Lewis uses the language (and concept) of possible worlds in place of "alternative possibilities."

[28]An examples adopted from Austin (1946) also discussed in Schaffer (2008).

[29]Many contextualists consider knowledge to be infallible within the context. That is, in order for a relevant possible world where $\neg p$ to be eliminated, it must be known that $\neg p$ with certainty. That is, "knowledge in light of possibilities of error sounds like a contradiction" Hendricks (2006, p. 67). However, for my purposes I will allow for the "evidence" that $\neg p$ to be fallible.

provide a rigorous technical definition (as the self-conscious informality of his formulation indicates!). What constitutes a possibility that one can properly ignore? This is left unclear. Some options are: significance, probability, relevance, resemblance, and salience.[30] This list is not exhaustive. Each has its own problems as a precise solution. For example, there is strong reason to doubt that just probability will work. The lottery paradox seems to show precisely that an alternative possibility's being extremely unlikely is not sufficient in isolation to justify ignoring it. Although Jane may be arbitrarily close to certain that she will lose the lottery, she may not properly ignore the alternative possibility where she happens to win.

Traditional attempts to specify the class of legitimately ignorable alternatives based on significance, relevance, world-resemblance, and salience share a common problem: namely, identifying just what determines whether an alternative possibility counts as relevant, salient, significant, and so on. As a basic necessary condition, Stine (1976) suggests that an alternative is relevant only if one has some reason to think that it's true. This seems intuitive: I know that that's a zebra because I don't have any reason to think that it's just a cleverly painted mule.[31] In fact, I have positive reasons to doubt that it's a cleverly painted mule because, perhaps, the zoo I'm attending is a reputable zoo and wouldn't risk its reputation. However, is something's being merely possible enough reason to think that it's true? This is left unclear. If something's merely being possible, or where one's evidence for $\neg p$ in subcases where $\neg p$ is not appreciably different in cases where $\neg p$ is false, then this suggests, pace Hawthorne, that all possibilities are relevant, and we arrive at the skeptical problem.

[30]Dougherty & Rysiew (2009) offer a "significance" version; Hawthorne (2004) offers a "salience" version.
[31]A reference to Dretske (1971).

One attendant problem for such views (*i.e.* salience, significance, etc.) is whether (mere) reflection or idle inquiry can too easily change the set of relevant alternatives that one must eliminate simply by bringing up a previously unconsidered further alternative. Suppose that I am asked whether there is a canary or raven in the garden. In this context I have evidence against the raven-in-the-garden world and (assuming that there is a canary in the garden) I may be said to know that there is a canary in the garden. But suppose that my interlocutor asks if it's possible that I'm dreaming. On some measure of relevance, this would make the I'm-dreaming world relevant to whether I know that there is a canary in the garden; I would now have to provide evidence against this possibility in order to eliminate it. If I could not, then I would not know that there is a canary in the garden, as I would not have eliminated all of the relevant alternative possibilities. But this seems to unduly privilege the power of an interlocutor to make relevant some unconsidered alternative. Moreover, this would suggest that cautious and reflective epistemic agents, who consider more alternative (and potentially less probable) possibilities, will know less than an incautious agent who considers fewer. The philosopher, after all, may doubt that she has hands, while a carpenter will not. Does the carpenter know that she has hands when the philosopher does not? We would not want this to be a consequence of a RAT.

It is a difficult task to explicate the conditions under which we can properly ignore alternative possibilities by deeming them irrelevant. Difficult tasks beset many theories, so I don't mean to suggest that we should rule out the possibility of an adequate account of what constitutes the necessary and sufficient conditions for properly ignoring alternative pos-

sibilities. At a minimum, each member of the RAT family is latching onto something that appears intuitively correct: namely, that there is *some* story worth telling about how one can properly ignore some alternative possibilities in coming to know a proposition. This is especially true if one holds both that one cannot know in the lottery case and that one can know in the ordinary case.

For my part, I suggest that an alternative possibility may be deemed irrelevant if its truth would imply the falsity of some, or enough of a particular class of background beliefs and/or theories about the world.[32] That is, the supposition of an alternative possibility would trigger a rational re-examination and restructuring of a significant class of background beliefs about the way the world works. I will refer to this as destabilization. For my present purposes, I will leave this concept deliberately vague. I will refer to this account as a *Non-Destabilizing Alternatives Theory*, or NDAT.

Ignoring possibilities based on their impact on background beliefs is not an entirely unfamiliar concept. The impact on background beliefs in NDAT is qualitative, rather than quantitative. We must avoid cashing out destabilization in terms of the number of background beliefs impacted. Suppose that Jane believes B1 and that she finds out that she actually won the lottery. It is natural to suppose that she would change a tremendous number of beliefs about her future (based on what she can do with the prize). But this is not

---

[32]Smith (2010) proposes something similar, but importantly different from my view. He argues that an alternative is relevant only if one's evidence *normically* supports the proposition. Evidence normically supports a proposition when the presence of the evidence but falsity of the proposition demands explanation; evidence that doesn't normically support a proposition doesn't have this property. Losing a lottery doesn't demand a special explanation, so my probabilistic evidence that my ticket will win does not normically support my losing the lottery. However, my seeing my cup and feeling its weight right before I close my eyes normically supports the proposition that I'm holding my cup. If I were to have those experiences and the proposition were false, this would demand a special explanation.

what I mean. Instead, destabilization refers to the change to more general beliefs and theories about how the world works.

There is a general theory of objects such that things such as cups cannot instantaneously (and for no apparent reason) become other objects like black holes. Science is concerned with articulating and compiling evidence for such theories. Suppose that Jane forms the belief that her cup will not suddenly become a black hole. Call this B4.

B4 My cup will not suddenly become a black hole.

Suppose that it's possible for B4 to be false. The falsity of B4 is therefore an alternative possibility. I suggest that if Jane were to find out that B4 is false, then this would be a destabilizing event. Her general theory of objects and nomological understanding of the world would be radically altered. But it's not the *number* as much as the *quality* of the background beliefs that would need to change on the supposition of B4 being false. My argument is that one may properly ignore destabilizing alternative possibilities in coming to know or assert some proposition.

In addition to destabilization involving a certain quality of impact to background beliefs rather than quantity, destabilization is a graded rather than a threshold concept. The supposition (or the coming to be) of some alternative possibilities will be more or less destabilizing than others. If I were to learn that cups could regularly suddenly become black holes, then this would be especially destabilizing to my background beliefs and theories about the world. However, if I were to find that my belief that my car is still where I parked it is false (since it has actually just been stolen), then this may be slightly destabilizing for my beliefs about

130

the security of the faculty parking lot, but it would be far less destabilizing than learning about cups becoming black holes. In the former, I would merely be rationally required to re-examine my beliefs about how quickly thieves can steal a car; in the latter, I would be rationally required to re-examine my views on the basic physical laws and probability assessments of whether objects can turn into singularities.

I suggest that a concept useful in explicating destabilization can be found in *consilience.* Generally speaking, consilience is a property of a body of beliefs, theories, or evidence, arising in distinct explanatory domains, which are mutually evidentially explanatorily. In the philosophy of science consilience has often been taken as an indicator of good theories, broadly on the grounds that cross-domain, inter-theoretic explanatory unity is a powerful element or indicator of correctness (on pretty much any notion of correctness).[33] For example, consider what has become known as the "modern synthesis" of (evolutionary) biology into what has also been referred to as "neo-Darwinism."[34] The synthesis of Darwin's theory of evolution by natural selection depended heavily on finding a vehicle for genetic heredity. When it was found (in the form of DNA) this led to successes in disparate fields such as economics (evolutionary game theory), psychology (evolutionary psychology), computer science (evolutionary programming and genetic algorithms), and many others. With each new success came more support for the synthesized theories of genetic heredity and evolution through

---

[33]Whewell (1840) was the first to introduce the concept. It has since been used in a number of contexts in the Philosophy of Science. However, I wish to distance my use of the concept from that of Wilson (1998). Consilience is sometimes manifested in a theory's producing successful predictions and explanations in areas where it was not expected to do work. For example, Newton's theory of gravity was not only able to explain the movement of celestial bodies, but also terrestrial motion. Moreover, it also both predicted and explained tidal motion through the gravitational influence of the moon.

[34]For a detailed description see Mayr (1991).

natural selection.

Now, since consilience pertains to hypotheses, theories, and beliefs, it is consistent with the possibility of disconfirming observations even for highly consilient theories. Some such observations would necessitate small modifications to a body of beliefs or hypotheses, but others would require a more radical restructuring of one's beliefs about *how the world works*. It is these latter events which I am describing as destabilizing. Irrespective of whether we ultimately choose to use the particular notion of consilience to underwrite the concept of destabilization, I think that it's plausible to see the intuitively appealing shared component of RATs as destabilization.[35]

Destabilizing events, though possible, are often ignored as being not worth consideration. I have not given a general definition of what counts as knowledge, though. I have merely been arguing for an account of when we can ignore (irrelevant) alternative possibilities in coming to know some proposition. Some might wonder what it means to knowledge if some destabilizing event, which was ignored as an alternative possibility, actually came about. If subject S ignored $\neg p$ because it is destabilizing, it does not follow that S knows $\neg p$ (this is a mistake that some RAT theorists have made). Suppose that I'm looking at what I take to be a zebra. I don't have any reason to think that it's a cleverly painted mule, and, in fact, I have good reason to doubt that it's a cleverly painted mule. If it were to turn out that what I'm looking at *is* a cleverly painted mule, then I would consider this destabilizing to my background beliefs about the zoo and about the passability of a cleverly painted mule for a genuine zebra. So I ignore the irrelevant possibility that what I'm looking at is a cleverly

---

[35]This is where my view dovetails with the view expressed in Smith (2010).

pained mule. However, I have only *ignored* the alternative possibility, I haven't *eliminated* it. When an agent eliminates an alternative possibility $\neg p$, the agent comes to know $\neg p$. However, when an agent merely ignores an alternative possibility $\neg p$, the agent does not come to know $\neg p$. Ignored alternatives do not thereby become known. So in ignoring the irrelevant possibility that what I'm looking at is a cleverly painted mule, I don't thereby know that it's not a cleverly painted mule.

The interplay between background beliefs, ongoing belief-formation, and observations of the world is an important feature of ordinary non-lottery cases such as A6 and fallible propositions in general. If John were to find that, when he asserted A6, someone had actually surreptitiously replaced his cup with a similar feeling object, this would more likely create a rational requirement to reform his background beliefs about the feasibility of making such a surreptitious switch quickly and without his noticing. In this case, if John were to find out that his belief about holding his coffee cup were false (because, in the instant between perceiving his cup and closing his eyes, someone replaced his cup with a cup-feeling non-cup), then this would be a destabilizing event. However, he wouldn't think that it was sufficiently analogous to the lottery case as though he just won the lottery to the I'm-not-holding-my-coffee-cup lottery.

This is consonant with the sense that there is something importantly different between lottery propositions like A1 and fallible propositions like A6. In ordinary cases our background beliefs and beliefs about underlying mechanisms are, or ought to be, importantly sensitive to outcomes. One of the reasons why lottery-talk is so useful, by contrast, is that

it abstracts away from questions about fairness, randomness, or underlying mechanisms. No matter how unlikely winning a given lottery may be, the prospect of winning requires no change to one's background beliefs about how the world works. Indeed, if Jane were to win her lottery and thereby reform her beliefs about the likelihood of winning a similar lottery, we would criticize her for committing a statistical fallacy. Suppose that Jane wins her 10,000 ticket lottery (and it is a fair lottery) and on the basis of this forms the belief that the odds of her winning must have been greater than 1:10,000. It would be appropriate to chastise her for committing a fallacy. Winning a large lottery may be unexpected, but someone has to win: why not her? Of course, winning an unlikely event may be evidence for an alternative explanation (such as a rigged draw), but in the lottery case this is insufficient to justify the belief. Thus, there seems to be something importantly different between the lottery and ordinary cases.

Now, one might object that this seeming distinction can be blurred by consideration of a brain-in-vat hypothesis.[36] For example, suppose that Bill is a BIV. The world is generally exactly the way that Bill thinks it is except that he is the result of an ambitious neuroscience experiment: he's the first viable BIV ever! The objection is that Bill, upon learning of his BIV status, wouldn't feel a rational requirement to change his more general background beliefs about how the world works. For example, he would still be able to make predictions about people's behaviour and the interaction of objects as well as he could before. He would just have to repair his local beliefs about his relationship to the world and people in it. The only difference his discovery makes is that being a BIV is now a live alternative possibility.

---

[36]My thanks to John Turri for raising this objection.

134

Therefore, Bill's learning that he is a BIV may not be a clear case of a destabilizing event. Since I have argued that we may ignore destabilizing alternative possibilities, this sort of BIV hypothesis would not be a proposition that Bill (or anyone else) could reasonably ignore as an irrelevant possibility prior to learning that he is a BIV. This would be contrary to the intuition that BIV hypotheses are irrelevant alternative possibilities. Therefore, the objection is that NDAT cannot explain why we ignore BIV-style alternatives since such possibilities may not be destabilizing.

The objection is not compelling, however. It asks us to consider Bill's epistemic state upon learning that he is a BIV even though the world is generally how he thinks that it is. If, prior to learning that he is a BIV, Bill thinks that there are tables and chairs, that the Toronto Maple Leafs are unlikely to win the Stanley Cup this year, and that it snows in Canada, then all of these propositions are still true. However, his beliefs about his relationship to the world are all false. If this objection were successful, then we would expect that learning that he is a BIV, even though the world is generally the way that he thinks it is, would not trigger a restructuring of Bill's background beliefs about the world. Bill could conceivably still go about placing wagers on the Maple Leafs not winning the Stanley Cup this year and do so rationally. After all, his beliefs about the quality of the team are just as true as before he learned about being a BIV. All that has changed is that he no longer actually has hands or the physical relationship to the world that he thought he had prior to learning about being a BIV.

Since I have argued that the impact on beliefs is qualitative rather than quantitative,

perhaps such a case would not be as destabilizing as traditional BIV hypotheses (where the world is not generally as we think it is). I have argued that destabilization is, after all, a graded concept: some events may be more or less destabilizing than others. I suspect that the intuition surrounding learning that one is a genuine BIV is that such circumstances would be strongly destabilizing; though, admittedly, perhaps less so than more traditional BIV/evil-demon hypotheses.[37] The sorts of propositions that one forms is heavily influenced by one's beliefs about one's relationship to the world. If Bill were to learn that he is a BIV, then it's hard to predict how he would react, but our popular culture apparently shares the intuition that he wouldn't take it very well: it would be highly destabilizing.[38] Bill's preferences for watching the Maple Leafs each year may be predicated on the benefits of winning wagers, such as being able to afford better tickets for the next game. But if he is a BIV, then there's no benefit to winning: the neuroscientists could merely provide a different input to give Bill the sensation of being in the best seats in the arena. It's plausible to consider many of one's life plans as strongly connected to one's beliefs about one's connection to the world. So learning that one is a BIV, even if the world is generally the way that one thinks it is, would be strongly destabilizing. I suspect that one would change many of one's life plans upon learning that one is a BIV.

---

[37]I say "genuine" since we're meant to imagine Bill being a brain in a vat rather than something less amazing such as learning that brains in bodies are, in some sense, brains in vats.

[38]Consider Movies such as *The Matrix* (1999), *Possible Worlds* (2000), and *Vanilla Sky* (2001).

5.6   Re-Examining Hawthorne

This notion of an asymmetry between lottery propositions and fallible propositions, along the lines of destabilization, can be more fully elucidated by discussing one of Hawthorne's examples (taken from Vogel): the Heartbreaker hole-in-one case.[39] I will simply refer to this example as the Heartbreaker. Sixty (amateur) golfers in a tournament must all play a short but difficult par 3 known as the 'Heartbreaker' hole. Holes-in-ones in golf are very rare even for touring professionals. For professionals the odds of a hole-in-one are roughly 1 in 3000; for an amateur roughly 1 in 12,000.[40] Based on this, prior to the tournament it would be reasonable to form the belief B5.

B5: Not all sixty golfers will get a hole-in-one on the Heartbreaker.

Roughly, the odds of all sixty golfers getting a hole-in-one on the Heartbreaker (*i.e.*, B5 being false) are astronomical: 1 in 5.63×10^{244}. This number is larger than all of the nanoseconds in the universe since the Big Bang by many orders of magnitude. Hawthorne argues that the appropriate response to finding oneself in the situation where all sixty golfers make a hole-in-one is the same as finding oneself in the situation where one wins an extremely large lottery. This would suggest that no restructuring of one's background beliefs would be required because no such response to winning a lottery is required (or perhaps even permitted). However, I think that this does not accurately capture our intuitive responses to such cases as the Heartbreaker.

---

[39]Hawthorne (2004, pp. 12 - 13 and 17 - 20); Vogel (1999, p. 165).
[40]Kindred (1999). I have rounded down for simplicity.

Consider what it would really be like to witness the event of all sixty golfers making a hole-in-one. The intuition, I think, seems to be that after, say, 8 or 9 holes-in-one, the next hole-in-one would cease to be surprising. Why? Partly because such an event is astronomically unlikely, we'd search for a more probable explanation; namely, that the game is rigged to result in holes-in-one on the Heartbreaker. One of the structural presuppositions of being able to interpret the ordinary cases as lotteries is that they're fair: that there's no foul play. Additionally, the possibility space must be divided such that one doesn't have any appreciably better reasons for thinking that one subcase obtains rather than another. However, partly based on background beliefs about the implausibility of all sixty golfers honestly making hole-in-ones, we ignore this possibility and search for alternative explanations, such as cheating, or a local topography that funnels shots into the cup. Even if no such better explanation were immediately forthcoming, the intuition would still be to decline believing that all sixty golfers achieve their feats under the sort of conditions we initially took to obtain when we first evaluated B5. We intuitively ignore that possibility and assume that some conditions or other obtain under which the supposition of the falsity of B5 would not be destabilizing. We would first attempt to exhaust all better explanations before being driven to the belief that each golfer honestly makes a hole-in-one.

However, if we suppose that they are honestly making these hole-in-ones, I suggest, we would intuitively be inclined to re-examine our background beliefs about the likelihood of such an event occurring. This would be destabilizing. However, contra Hawthorne, if the Heartbreaker case can be properly reconstructed as a lottery, then such an inclination would

be absent. The presence of the inclination to re-examine our background beliefs and theories about the world (given that all 60 golfers honestly made their hole-in-ones) suggests that the Heartbreaker case cannot be properly reconstructed as a (very large) lottery. There seems to be an intuition that one would not be committing a statistical fallacy by re-examining the likelihood of the Heartbreaker event given its occurrence (because we re-examine our background beliefs about the world). Therefore, this indicates an important asymmetry between lottery and fallible propositions. In fact, if Jane were to repeatedly win large lotteries, then this would constitute stronger evidence of cheating (much like, I have argued, the Heartbreaker case) and a stronger candidate for destabilization.

For Hawthorne, mere attention to the possibility of error is not sufficient to vitiate knowledge. Instead, knowledge is vitiated by attention to an epistemic possibility of error. He calls this the Epistemic Possibility Constraint (EPC).[41]

> EPC: If it is epistemically possible for S that $\neg p$ then S does not know that p.

However, in order for us to take the skeptic seriously, the skeptical alternative possibilities of fallible propositions must be epistemically possible. If John is holding his cup and is asked whether he can know A6 because, after all, it's possible that his cup has been surreptitiously switched, John will entertain such an alternative possibility as epistemically possible. In fact, it's not clear what difference Hawthorne has in mind between the mere attending to a possibility and something being epistemically possible. Typically, something is epistemically possible for an agent if some proposition (or state of affairs) is consistent with what one

---

[41]Hawthorne (2004, pp. 24 - 28), and elsewhere.

knows.[42] All of the alternative possibilities discussed count as epistemically possible in this sense: whether one is a brain-in-a-vat, or that a coffee cup has been surreptitiously replaced with a non-cup, whether a table has temporarily changed due to a massive confluence of quantum effects, or whether all 60 golfers honestly make a hole-in-one on the Heartbreaker.

As indicated above, Hawthorne argues that if it is epistemically possible that $\neg p$ then one cannot know that p (EPC). But this is just what it means to be a fallibilist: that one may know p despite $\neg p$ being consistent with one's evidence.[43] Fortunately, since most people favour fallibilism, we can simply reject EPC and I have given grounds for doing so. Consider John entertaining the possibility that A6 is false because someone has switched his cup. Naturally, he will ignore the possibility. He recognizes and accepts it as an epistemic possibility, but he dismisses it as irrelevant. If A6 is true, and he has good reason to think that it is based on his perceptual experience of the cup and his background theory of objects and how the world works, then John knows (and may warrantedly assert) A6. His entertaining that A6 being false as epistemically possible (because John is a fallibilist too!) does not vitiate his knowledge.

Such is the case for fallible propositions when epistemic agents entertain skeptical alternative possibilities: they are epistemically possible but irrelevant. Why? Because they're destabilizing. One cannot appeal to their improbability, for this is the lesson of the lottery paradox. Therefore, the case is not the same for lottery propositions. Why? Because they're not destabilizing. We do not have grounds to ignore the alternative possibilities where $\neg p$.

---

[42]See DeRose (1991), for example.

[43]One could argue that a strong motivating factor behind Fantl & McGrath (2009) is taking issue with the tension between being a fallibilist but subscribing to EPC.

But, importantly, for both lottery and fallible propositions the alternative possibilities where $\neg p$ are considered as epistemic possibilities. It's just that in the latter, but not the former, such possibilities can be properly ignored and therefore do not vitiate knowledge or assertibility.

NDAT also performs better than salience versions of RATs, such as Hawthorne's, sharing their virtues while avoiding two problems common to such salience accounts. The first of these problems is explaining which alternatives count as salient given the context. Is the (epistemic) possibility that John's cup could have been surreptitiously replaced a salient one? What about whether my car has been stolen from the parking lot? The second problem is the implausible implication that cautious epistemic agents will ipso facto know less than incautious agents. The epistemologist may habitually entertain a larger array of alternative possibilities than an incautious layperson. Or an irritating interlocutor may cause an agent to lack knowledge merely by increasing the range of seemingly salient options by mentioning such possibilities as brains-in-vats and quantum-shift coffee cups. But NDAT avoids this problem. Whether an alternative possibility counts as salient on some reckoning is irrelevant to whether the alternative can be rationally ignored or must be eliminated. Rather, what matters is whether the supposition of the alternative would be destabilizing.

5.7   Implications for Norms of Assertion

In this chapter I have attempted to establish an asymmetry between the epistemic properties of lottery propositions and fallible propositions. One can know fallible propositions

such as that I am writing at a computer, but cannot know lottery propositions such as that my ticket will not win in tomorrow's lottery. In order to know some proposition one must ignore irrelevant alternative possibilities and eliminate, or be in a position to eliminate, the relevant alternative possibilities. Although specifying precisely what constitutes grounds for ignoring irrelevant alternative possibilities has been notoriously difficult, I have argued that an appealing account of the asymmetry can be found in a relative alternatives approach that builds on the differences between destabilizing and non-destabilizing observations, which I have called NDAT. If the supposition of an alternative possibility would trigger a significant restructuring of our background beliefs about how the world works, then such an event is destabilizing. Discovering that, contrary to every sustained indication, I was not writing at a computer would surely destabilize much of my understanding of the world and the reliability and meaning of my experiences across the board. Discovering that my ticket had won the lottery would be a great surprise, but would require no restructuring of my thoughts on how the world works. Indeed, in doing so I would commit a statistical fallacy. Destabilizing alternative possibilities, though possible, may be properly ignored in the first instance. Non-destabilizing alternative possibilities (such as one's actually winning a lottery) cannot be ignored. In the case of fallible propositions, but not lottery propositions, the alternative possibilities which the skeptic uses as a wedge may be ignored because they're destabilizing. Provided that one also eliminates the relevant (*viz.*, non-destabilizing) alternative possibilities, then one can know such propositions.

In this way, NDAT retains Williamson's basic intuition INT1, explaining why an agent

fails to know in the lottery case. Furthermore, unlike Hawthorne's, NDAT is able to retain Williamson's second basic intuition INT2. It explains how one can know in the ordinary case. However, NDAT does not depend on identifying the pathology of asserting lottery propositions (such as A1) with the "merely probabilistic" nature of one's evidence – a move that has problematic implications for Williamson's own account. By rejecting the insufficiency principle, the NDAT version of a Relevant Alternatives approach can remain fallibilist and yet reject knowledge in the lottery case, while allowing for knowledge in ordinary cases.

Although we can now effect a separation between lottery and fallible propositions, such that the latter but not the former are knowable, this is not yet to comment on the implications for norms of assertion. The lottery paradox is one of the principal pillars in the case for KNA: what best explains the unassertibility of (A) and statements like it is that one cannot know that a given ticket will lose. This doesn't seem particularly well explained by the truth norm, for example, as discussed in Chapter 4. Lottery propositions are, by definition, arbitrarily well supported by one's evidence and yet remain intuitively unassertible even when the proposition asserted happens to be true. I discussed Weiner's attempt to invoke a Gricean explanation for the unassertibility of lottery propositions. To recapitulate, Weiner argues that TNA can explain the unassertibility of lottery propositions on the grounds that doing so breaks basic Gricean norms of conversation not specific to assertion. So although one doesn't break the norm of assertion – the truth norm – it violates basic conversational norms, and this explains the unassertibility.

This explanation may be fine, as far as it goes, but there still seem to be cases of unassert-

ible lottery propositions which don't admit of a Gricean explanation. Suppose that Jim buys a ticket to a lottery with a large payoff but a very large number of tickets. Sally asks him if he bought a ticket, and he responds, "I bought a ticket, but it won't win." The second conjunct is an assertion of a clear lottery proposition, which seems intuitively unassertible. A Gricean explanation is not available: Jim is responding to Sally's question. Her question seems to presuppose that buying a ticket might make sense, so Jim's assertion that his ticket won't win doesn't violate Gricean norms in the way that Weiner argues Sarah's assertion to Alice does. So the TNA advocate needs a different explanation for the unassertibility of this sort of lottery proposition.

Fortunately I have given the TNA advocate a solution in this chapter: speakers cannot satisfy Weiner's requirement that speakers have evidence for the truth of their assertions, perhaps in terms of secondary propriety.[44] I've argued that such justification is never present for lottery propositions, only for fallible propositions. So the TNA advocate can account for the lottery paradox linguistic data by appealing to arguments against speakers having adequate evidence to warrantedly assert.

The principal problem with this response, though, is that it seems to better support a reasons-based norm than either KNA or TNA. Although KNA and TNA can both explain the impropriety of asserting lottery propositions by appealing to the lack of (epistemic) justification for such assertions, that the central feature missing in such assertions is adequate (epistemic) justification seems most directly to support an assertoric norm that places such a

_____

[44]However, the TNA advocate may now be hanging their case on the plausibility of primary/secondary propriety distinction.

requirement at its core. That is, a reasons-based norm, such as what I argue for in Chapter 8, most easily explains the data because the feature that renders such assertions unwarranted is the central feature of a reasons-based norm; namely, a lack of sufficient (epistemic) reasons! So while KNA advocates have used the lottery paradox as a pillar of support for KNA, we can now see how a variety of norms, including TNA and reasons-based norms, can also easily explain the data. And, furthermore, since KNA's explanation for why a speaker can't know a lottery proposition depends on a lack of epistemic justification, the data doesn't uniquely support KNA rather than a reasons-based norm.

In Chapter 6 I take up another pillar of support for KNA, Moore's Paradox, and again argue that the data is well explained by a reasons-based norm and, therefore, doesn't uniquely support KNA. Chapters 5 through 7, therefore, constitute a re-examination of some of the principal arguments made in favour of KNA which, I argue, actually support a reasons-based norm. And in each instance, I argue, the explanation of the linguistic data appeals to a speaker's lack of sufficient epistemic reasons for warranted assertion. Warranted assertibility seems most closely associated with having adequate reasons than an assertion expressing either knowledge or truth.

Chapter 6

Moore's Paradox and Assertion

6.1 Introduction

Moore's Paradox looms large in current debates over the norms of assertion. It is generally uncontroversial that assertions such as, "It's raining, but I don't believe that it's raining," are, as Moore put it, "absurd." Call such assertions, and their attendant absurdity, *Moorean*. When it comes to evaluating competing proposals regarding the norms of assertion, it is widely held that adequately explaining the absurdity of Moorean assertions is a virtue. It has also been alleged by advocates of various norms of assertion – in particular, by advocates of the knowledge norm – that their view is especially well-positioned to explain Moorean absurdity. In many cases, Moore's Paradox has been a central battle ground between competing proposals for the central norm of assertion. In this chapter, I will argue that several competing norms of assertion can well explain Moorean absurdity, in which case none gains a comparative advantage. Consequently, the capacity of a norm to explain Moorean absurdity

is better thought of as an adequacy condition for a proposed norm of assertion.

In this chapter I will review some influential arguments involving Moore's Paradox in its various forms, with a particular emphasis on their roles in the debate over assertoric norms. In order to properly understand the explanatory virtues alleged to accrue to a proposed norm of assertion that is capable of explaining MP cases, we ought first to understand the source of the pathology of Moorean statements. I thus begin in Section 6.2 by discussing Moore's original formulation of the paradox and his attempt at solving the problem. In section 6.3 I briefly discuss a number of different forms that have been labelled "Moore's Paradox" and have been used in arguments on assertoric norms. In exploring the pathology of Moorean statements in Section 6.4, I discuss Sorensen's taxonomy of Moorean cases, and consider some reasons why we may doubt the centrality or relevance of some of the proposed Moorean forms listed in Section 6.3. I will then canvass, in Section 6.5, the various arguments that have been given for how the principal candidate norms of assertion can adequately explain the oddity of the important forms of MP. Finally, in Section 6.6 I will argue that each of the candidate norms is able to adequately explain all of the relevant MP data and that we should therefore not treat MP as providing evidence in favour of any given norm. We should instead treat a candidate norm's ability to explain MP as an adequacy condition.

6.2 Moore and Moore's Paradox

The original formulation of Moore's Paradox is G.E. Moore's own "I believe it's raining, but it isn't."[1]

    (1) I believe it's raining, but it isn't.

This is known as the *commissive* form of the paradox, which is often represented with the general form "p, but I believe that $\neg p$." It's commissive since the second conjunct involves a positive claim about belief: the speaker claims to believe that *not-p*. We can contrast this with the *ommissive* form, as in (2):

    (2) It's raining, but I don't believe that it's raining.

It's ommissive since the second conjunct doesn't involve a belief about p (or *not-p*) per se; rather, the speaker claims not to believe that p – it doesn't follow that the speaker believes that *not-p*, which is the case in the commissive form.

    Although (1) and (2) might at first appear to be logical contradictions, they are not. A contradiction is a necessarily false proposition, but both (1) and (2) may be true. It's possible for it to be raining and for a speaker to fail to believe that it's raining. After all, it's generally uncontroversial that agents can have false beliefs or even false meta-beliefs. The paradox, then, seems to arise when agents either assert or occurrently believe a Moorean proposition.[2] The difficult task is explicating from whence the pathology arises and why it only seems to arise when an agent believes or asserts such sentences.

---

[1]Moore (1962, p. 291). Though, we first find the paradox in Moore (1942). For a thorough discussion of Moore's Paradox, see Williams & Green (2007).

[2]Sorensen (1988) distinguishes between Moorean absurd *propositions* and *sentences*.

While (1) and (2) are not, strictly speaking, contradictions, one still might think that their assertion straightforwardly entails a contradiction. That is, one who asserts (2) *appears* to assert that one both believes that it is raining and that it is not raining.[3] In doxastic logic we would represent this as Bp & ¬Bp. But it's not clear that this is the correct interpretation. What is asserted is in fact an instance of p & ¬Bp. While one may argue that believing or asserting this is absurd by the lights of epistemic constraints on rational belief, there is no contradiction directly entailed by the assertion (or the belief).[4] Some fairly powerful principles of assertoric and doxastic logic would be required to generate a logical contradiction from the mere assertion itself.

Moore's own analysis of the oddity of such assertions stems from what he views as an implied inconsistency between what one asserts and how one represents oneself in making the assertion. Moore held that in asserting *p*, one *represents oneself* as believing *p*.[5] In fact, one may go further and say that in asserting *p*, one represents oneself as knowing *p*. By asserting "It is raining" as the first conjunct, the speaker represents herself as believing (or perhaps knowing) that it is raining. On this view, by also asserting "I don't believe that it is raining," the second conjunct seems to contradict what the speaker represents herself as believing by asserting the first conjunct.[6]

---

[3]However, it's significant that the sentences are in the form of "p & ¬Bp."

[4]See, for example, de Almeida (2001).

[5]Williamson (2000, p. 252) reports that this may have begun with Max Black: "In order to use the English language correctly, one has to learn that to pronounce the sentence 'Oysters are edible' in a certain tone of voice is to *represent oneself* as knowing, or believing, or at least as not disbelieving what is being said" in Black (1952, p. 31). C.f. DeRose (2002) and Kvanvig (2009).

[6]If we assume that belief distributes over conjunction, some have said that holding an ommissive form of MP, the claim not to believe that *p* "falsifies" the claim that *p*. This is different from the claim that there's a contradiction in holding a Moorean absurd belief. See, *e.g.*, Williams (1994), Littlejohn (2010). Compare de Almeida (2001). I discuss this below.

Moore's analysis of the source of the absurdity of Moorean propositions thus comes from what is strongly insinuated by one's assertion. I say "insinuated" to avoid the logical meanings of "entail," "imply," and implicate. The meaning of "insinuate" is to hint or suggest without logically implying, entailing, or implicating. However, Moore does use "imply." So although (2) is not a contradiction, and its assertion does not entail a contradiction, its assertion strongly insinuates a contradiction.

Moore considered a second set of propositions which seem to share the same degree of oddity as (1) and (2). These have become called the knowledge version of MP, as in (3):

(3) Dogs bark, but I don't know that they do.[7]

The explication of knowledge versions such as (3) is the same as for the belief versions: "By asserting p positively you imply, though you don't assert, that you know that p."[8] But just as with the belief versions of MP, the proposition asserted is not a contradiction, nor does it entail a contradiction. If asserting p means that a speaker represents herself as knowing p, then asserting the second conjunct is inconsistent with what is represented by asserting the first conjunct, but not in the sense of being, or entailing, a contradiction.

6.3 More Moorean Propositions

Michael Huemer offers a list of various formulations of Moorean propositions which can be found throughout the literature.[9]

---

[7]Moore (1962, p. 277). See also Unger (1975) and Stanley (2008).

[8]Moore (1962, p. 277). It's important to note that Moore does not mean "imply" in the logical sense of entailment, but rather in the sense in which I use "insinuate."

[9]Huemer (2007).

(4) It is raining, but I do not believe that it is.

(5) It is raining, but I believe that it is not.

(6) It is raining, but I do not know that it is.

(7) It is raining, but that isn't true.

(8) It is raining, but I have no justification for thinking so.

(9) It is raining, but my reason for thinking so is false.

(10) It is raining, but there are (non-misleading) facts that neutralize my reasons for believing that.

(11) It is raining, but my belief that it is was formed in an unreliable way.

(12) It is raining, but I would believe that even if it were false.

(13) It is raining, but I am not sure that it is.

(14) It is raining, but it is not certain that it is.

From deAlmeida, we may also include the following:

(15) It is raining, but I doubt that it is.[10]

There is a question whether all of these are, properly speaking, Moorean propositions. In Section 6.4 I turn to Sorensen's taxonomy and definition of a Moorean proposition. His definition, in addition to Jason Stanley's subsequent analysis, will give us reason for thinking that propositions such as (13) and (14) are not, properly speaking, Moorean.[11]

---

[10]de Almeida (2001).
[11]For my discussion of Stanley's analysis, see p.162 below.

6.4   Sorensen's Analysis and Taxonomy

Roy Sorensen provides one of the most detailed analyses of Moore's Paradox which includes a detailed taxonomy and definition of Moorean propositions along with an attempted solution to the paradox.[12] He divides instances of Moorean propositions into the categories of "pure" and "impure" and further divides them into obvious and non-obvious cases. A paradigm example of an obvious Moorean proposition is (4). We can easily identify that a speaker who asserts or believes (4) asserts something obviously odd. But there are "non-obvious" Moorean propositions which are not so easy to diagnose, such as (16), spoken by a male.

>   (16)   The atheism of my mother's nieceless brother's only nephew angers God.

This is Moorean since the subject of the proposition is the speaker himself. So (16) is equivalent to "My atheism angers God" which is equivalent to "God exists, but I don't believe that God exists." The latter is much more clearly Moorean. In cases such as (16) we can, as it were, perform a "Moorean transformation" on the proposition to more clearly see its Moorean character. But we are not yet much closer to understanding what constitutively defines Moorean propositions.

Both (4) and (16) constitute pure Moorean propositions in Sorensen's terminology. A proposition is a pure Moorean proposition if it can only be diagnosed as defective by appealing to its Moorean character. By contrast, an *impure* Moorean proposition is one that can be diagnosed as defective without appealing to its Moorean character. In understanding pure

---

[12]One motivation for his spade work is that it's in service of his argument that there are some true but unknowable propositions which he calls *blindspots*. For example, "No one knows that p, yet p." Sorensen (1988, p. 52).

Moorean propositions, contrast (16) with (17) below:

(17) All bachelors are males, but I believe some bachelors are not males.

(17) shares the same structure of the typical "p, but I don't believe that p" of pure Moorean propositions, but (17) is different because the second conjunct clearly expresses a contradiction. So although we may criticize the assertion or belief of (17) on Moorean grounds, we more directly criticize it for including an obvious contradiction. In this sense (17) is an impure Moorean proposition.

Although the issue of whether Moorean propositions actually represent inconsistencies was briefly discussed in Section 6.2, Sorensen provides a much more nuanced treatment of the various types of inconsistencies we find in propositions and the way Moorean propositions should be understood as involving an inconsistency. Sorensen distinguishes between *direct*, *indirect*, and *patent* forms of inconsistencies, none of which he considers genuine cases of Moore's Paradox. A direct inconsistency involves a number of beliefs that express a contradiction; an indirect inconsistency is a case where we can derive a direct inconsistency; and a patent inconsistency is a case where the content of a belief or assertion is itself a contradiction. I will discuss them in turn.

Sorensen suggests that there are multiple forms of direct inconsistencies, such as the following.

(18) I believe that it is raining and I believe that it is not raining.

(19) I believe that it is raining and I don't believe that it is raining.

153

We can represent (18) as "Bp & B¬p," and (19) as "Bp & ¬Bp." He argues that they are inconsistent beliefs "because the inconsistent beliefs are directly opposed to one another across the same proposition: (Bp & B¬p)."[13]

Sorensen's point warrants discussion. In his discussion of Moore's Paradox he explicitly eschews relying on doxastic logic. In order for (18)–(19) to constitute direct inconsistencies they must be real contradictions: they must be necessarily false. So in order for Bp & B¬p to constitute a contradiction, it should be impossible for an agent to both believe that it is raining and believe that it is not raining. But this is not the case: all sorts of people have all sorts of beliefs which share the form of (18). DeAlmeida provides an example. Suppose that I visit "my psychoanalyst, whom I know to be an extremely reliable source of information about both my beliefs and the facts of my life, [and she] tells me: 'Your father loves you, but you believe he doesn't.'"[14] On this basis I may rationally come to believe "I believe that my father loves me, but I also believe he doesn't" which would have the same form as (18). Since I can, in this case rationally, believe this, Bp & B¬p is not necessarily false, and so isn't really a contradiction. Although the propositions believed cannot both be true, that I believe both of them can be true.[15] So something more is required to show how such beliefs are inconsistent, a point to which I return shortly. However, Bp & ¬Bp *is* a contradiction: I can't both believe that p and not believe that p.

---

[13]Sorensen (1988, p. 26).

[14]de Almeida (2001, p. 43). While this is not deAlmeida's purpose for the example, it serves my purposes well.

[15]However, the knowledge version – Kp &¬p – expresses a contradiction due to the factivity of knowledge.

The problem here with Sorensen's taxonomy is that there is an ambiguity in what it means for an agent to have inconsistent beliefs. This could mean that the propositions expressed in two beliefs are mutually exclusive. That is, both propositions cannot be true. The propositions expressed by the beliefs in (18) – "it is raining" and "it is not raining" – cannot both be true: it's either raining or it isn't.[16] In this sense, then, an agent could hold inconsistent beliefs but it doesn't follow that she thereby believe a contradiction. This would require the doxastic logical principle that belief aggregates under conjunction (*i.e.*, Bp & B¬p → B(p & ¬p)), which Sorensen eschews.

Alternatively, "inconsistent beliefs" could be mutually exclusive beliefs where my believing one proposition seems to falsify my claim to believe in another. In (19) my claim to believe that it is raining seems to falsify my claim to not believe that it is raining.[17] This "falsify" terminology is common throughout the literature. In one sense this seems obviously true: if I believe that it is raining, then it's false that I don't believe that it is raining. Suppose that I assert "I know that it is raining, but I don't believe that it is." Clayton Littlejohn argues that in this case "we have a situation where [my] second-order belief is transparently falsified by a fact about [my] own mind."[18] We can run this in two ways. First, if I really do know that it is raining, then since belief is a necessary condition for knowledge, I must believe that it is raining. Littlejohn thus suggests that my believing that it is raining makes false the second order belief that I don't believe that it is raining. Second, if I really don't

---

[16]My discussion avoids all discussion of para-consistent and non-standard logics where there may be true contradictions.

[17]Perhaps more importantly, the placement of the negation makes it irrelevant that belief is involved in the claims at all. It's just an instance of asserting "*p* and ¬*p*." So the connection with knowledge here might not be the point.

[18]Littlejohn (2010, p. 22).

believe that it is raining, then I cannot know that it is raining because believing that it is raining is a necessary condition for knowledge.

As a third example, we may interpret "inconsistent beliefs" as applying to an agent's holding beliefs that produce inconsistent rational *commitments*, and perhaps even commitment to a contradiction. This is eventually Sorensen's position: Moore's Paradox involves beliefs (and assertions) that rationally commit the agent to inconsistent beliefs: that is, beliefs involving mutually exclusive propositions.[19] Having a rational commitment to some proposition entails being required to provide reasons for believing – that is, endorse – some proposition. So being committed to endorsing a contradiction is problematic because one cannot simultaneously provide good reasons for believing p while also providing good reasons to believe ¬p.[20] In this sense, the beliefs are not inconsistent *viz.* expressing the belief of a contradiction such as B(p & ¬p), but instead they express an agent's rational commitment to holding contradictory beliefs such as (Bp & B¬p). (18) rationally commits an agent to endorsing mutually exclusive propositions (in the form of beliefs).

Continuing with Sorensen's taxonomy, we must also distinguish direct inconsistencies from indirect inconsistencies. One way to do this is by considering the Lottery Paradox (see Chapter 5 for a much more complete treatment). Suppose that Jane holds a ticket in a fair local lottery of 1000 tickets. She reasons, based on the low odds of her ticket winning, that "My ticket (ticket #1) will lose." By parity of reasoning she may form the same attitude to any other ticket such that she can construct the belief that "Ticket 1 will lose, and ticket 2

---

[19]I discuss what Sorensen means by rational commitment in more detail beginning on p. 159.

[20]Brandom (1998) is a particularly salient example of someone who advocates this view.

will lose...and ticket 1000 will lose." We're tempted to think that she's justified in believing each conjunct and, through a minimal principle of closure, she would be justified in believing the conjunction. (The paradox results because we think that she is not justified in believing the conjunction.) But since it is a fair lottery she is also justified in believing that at least one ticket will win. So this case potentially produces the inconsistent proposition (20):

(20) Although I believe of each ticket that it is not a winner, I believe that at least one of them is a winner.

Both propositions expressed in the beliefs in (20) cannot be true, so Sorensen argues that it expresses a contradiction. Indirectly inconsistent propositions do not take the forms of direct inconsistencies noted above for (18)–(19), but we may derive a direct inconsistency since what is expressed in the first conjunct of (20) logically implies the opposite of what is expressed in the second conjunct.

However, the same comments levied against the ambiguity in "inconsistency" and its cognates in Sorensen's analysis of direct inconsistencies also applies to his discussion of indirect inconsistencies. (20) is not really inconsistent or contradictory. Instead, it would only rationally commit an agent to inconsistent beliefs. One may argue, for example, that if Jane is justified in believing that one ticket must win (which should be granted in a fair lottery), then she can't be justified in her beliefs about each ticket not being a winner for *all* of the tickets. At best, as discussed in Chapter 5, she is justified in believing of any given ticket, but only for that ticket, that it will not win. She can't continue to form a justified belief for each and every ticket. That doing so leads to the rationally inconsistent (20) is the reason.

So while both conjuncts of (20) cannot be true, we can only derive a direct inconsistency with the form of (18) rather than (19): we can derive the form Bp & B¬p from (20). And since this form does not express a contradiction, at most (20) rationally commits the agent to inconsistent beliefs.

The final division in Sorensen's taxonomy of inconsistencies is distinguishing direct and indirect inconsistencies from patent inconsistencies, such as (21).

(21) I believe it is both raining and not raining.

We can represent patent inconsistencies as sharing the form B(p & ¬p). Whether patent inconsistencies are really distinct from direct inconsistencies is not a topic of interest for my purposes.[21] What is important is how Sorensen argues that Moorean propositions are neither directly, indirectly, or patently inconsistent. One primary difference between utterances of (1)–(2) and (19) is that in order to suppose that an utterance of (1) is inconsistent we must assume that speakers believe their assertions. While this is a common assumption in the literature, no such assumption is required to show that (19) seems inconsistent. The more obvious difference is that, as mentioned in Section 6.2, Moorean propositions can be, strictly speaking, true whereas inconsistent propositions cannot. A proper characterization of Moorean propositions will thus univocally explain pure and impure as well as obvious and non-obvious cases and distinguish their apparent unassertibility (*i.e.*, not warrantedly assertible) from other unassertible propositions such as "Pigs fly."[22]

---

[21]Because, for example, one might think that belief distributes such that we can represent (21) as having the form Bp & B¬p, which is a direct inconsistency in Sorensen's taxonomy.

[22]By that I mean that it's not warrantedly assertible since it's obviously false and lacks sufficient supporting evidence.

So what explains the oddity of Moorean propositions? Sorensen canvasses a number of proposals. One might naturally think that Moorean propositions are consistent propositions which are unassertible because the proposition is falsified on its assertion. "I am not speaking now," for example, seems falsified upon its assertion. However, Sorensen argues, this may be too narrow for a definition of the problem. Such a proposal may only work for non-obvious cases such as (16) but not obvious cases such as (1). Sorensen suggests that it's hard to find it plausible that one would assert something so plainly odd as (1), but we do find it plausible to construct non-obvious cases. Furthermore, (1) doesn't seem clearly made false by its assertion. So although there are some Moorean propositions which are falsified by their assertion, this isn't the case for all Moorean propositions.

A natural development of this first proposal is to define Moorean propositions as, strictly speaking, consistent but not warrantedly assertible. This is a feature apparently common to all Moorean propositions, but the definition would be too wide: it doesn't properly explain why "Pigs fly" is not warrantedly assertible. Let's assume that pigs can fly, but John has no evidence of this; in fact he has very strong evidence that pigs cannot fly. In this case "Pigs fly" would be a consistent but not warrantedly assertible proposition. But "Pigs fly" is not a Moorean proposition. So Sorensen argues that we should reject this proposal as too wide.

Sorensen argues for a treatment of Moore's Paradox, since largely adopted in the subsequent literature, that characterizes Moorean propositions as consistent propositions that *rationally commit* the speaker to inconsistent beliefs.[23] By commitment he does not require an agent to occurrently hold inconsistent beliefs, or even dispositionally believe them. In-

---

[23]C.f. Brandom (1983, 1998).

stead, we may understand commitment in deontic terms. Stress must therefore be placed on "rational" sense in which one is committed. One is not metaphysically required to actually hold the inconsistent belief to which an agent is committed. Suppose that Mary believes that James can fly. Since it's analytically true that if James can fly, then someone can fly, Mary is rationally committed to the belief that someone can fly.[24] If Mary were to believe that James can fly but disavow the belief that someone can fly, she would be making a mistake. James Pryor characterizes rational (doxastic) commitment as follows.

> Take a belief the subject happens to have, *e.g.*, his belief in P. Consider what would be the epistemic effects of his having (decisive) justification for that belief.[...]If one of the effects is that the subject has decisive justification to believe Q, then his belief in P counts as rationally committing him to the belief in Q – regardless of whether he really does have any justification to believe P.[25]

Sorensen argues that one ought not believe a proposition, even if it is consistent, unless it is rationally *completeable.* That is, one ought not believe a proposition if it commits one to a direct inconsistency. However, he does not require that agents work through the rational commitments of their various beliefs before coming to believe a proposition: that would be unreasonable as a general requirement. "Since logical truths are consequences of any proposition, absolute thoroughness (given that we believed something) would entail logical omniscience" and this is widely rejected.[26] So we do not require epistemic agents to believe, even dispositionally, all of the logical consequences of their beliefs. However,

---

[24]From Pryor (2004).
[25]Pryor (2004, p. 364), emphasis in the original.
[26]Sorensen (1988, p. 36).

Sorensen does argue that rational completeability is a norm of belief, albeit unattainable and a waste of cognitive resources to even attempt to approximate.[27] This is an extremely restrictive doxastic policy since it would be surprising to find a single person who conforms to this norm. Sorensen recognizes this and suggests that we should adopt the policy that it is better to commit an omissive error rather than commit a commissive error. He claims that "[t]he amount false beliefs harm us exceeds the amount true beliefs benefit us. Better therefore to abstain from belief than to suffer an inevitable net loss."[28] Thus, he suggests, "ommissive error is less bad than commissive error."[29] So he endorses the policy that we ought to maximize true beliefs while minimizing the risk of false beliefs.

Sorensen thus diagnoses the problem with Moorean propositions in terms of speakers (or believers) taking on rational commitments to inconsistent beliefs – a diagnosis with which I agree in broad outlines. Yet his own solution to the paradox is too extreme. Sorensen holds that commissive error is worse than ommissive error: that it is better to suspend belief than risk acquiring false beliefs. While knowing whether a prospective belief is rationally completeable is an unrealistic requirement, it's better to avoid believing something true than believing something that rationally commits one to an inconsistency. That is, it's best to be conservative with one's beliefs. Therefore, one should not believe Moorean propositions because of the risk of commissive error: it's better to avoid believing something true than to believe something false, or something that rationally commits one to an inconsistency.

---

[27]This might make us question the value of the norm. This objection is related to the discussions in Chapters 3 and 8.

[28]Sorensen (1988, p. 37).

[29]Sorensen (1988, p. 38).

However, such a position is not required in order to properly address Moore's Paradox. Moreover, the norm of completeability seems too strong, particularly given that he admits that it is unattainable and not worth attempting even to approximate. Rather, Sorensen already provides us with adequate tools to solve the problem: Moorean propositions are problematic insofar as they rationally commit the agent to inconsistent beliefs. Of course, some inconsistencies are permissible, particularly given a sufficiently large class of beliefs. The preface paradox is a common example: although I may be sufficiently justified for each of my individual beliefs, I am also justified in believing that at least one of my beliefs is false, even though I don't have evidence for which one(s) might be false. There is nothing pathological here. However, there seems to be something pathological about how transparently Moorean propositions commit speakers to inconsistencies.

The upshot of Sorensen's arguments that will carry throughout this dissertation is his characterization of Moore's Paradox as the assertion of (potentially) consistent propositions which rationally commit the speaker to inconsistent beliefs. This characterization now places us in a better position to revisit Huemer's list from Section 6.3. (4)–(7) are properly Moorean; on Sorensen's taxonomy they are pure Moorean propositions. The problem with Huemer's list is that (8)–(15) may not rationally commit speakers to inconsistent beliefs. As Jason Stanley has noted, there is nothing pathological about (13) and (14).[30] He notices that there is an important asymmetry between assertions such as "John is here, and I know he's here" and "John is here, and I'm certain he's here." We think, he argues, that the second conjunct in the latter assertion, but not the former, adds information not present in the first conjunct.

[30]Stanley (2008). Compare Unger (1975).

162

That is, adding "and I'm certain" isn't redundant in the way that adding "and I know" is. One could interpret Stanley's argument as implying that Moorean propositions involve the improper denial of a redundant second conjunct. That is, "and I know p" is a redundant second conjunct in "p and I know p." So "and I don't know p" is the denial of a redundant second conjunct in "p and I don't know p." The latter thus strikes us as contradictory because asserting "p" implies (intimates) that the speaker knows that p; so denying "and I know p" strikes us as odd. But since "and I'm certain that p" isn't redundant to an assertion of "p", denying that one is certain, while asserting p, is not odd in a Moorean sense.

Stanley nevertheless argues that propositions such as (13) and (14) are, properly speaking, Moorean. However, given Sorensen's characterization, it's not clear on what basis this could be argued. (13) and (14) may sound odd, but this is not sufficient to count as Moorean. As DeRose has suggested, "I don't have a special feeling for inconsistencies; I can sense some kind of clash, but cannot distinguish my sensing of an inconsistency from my sensing of whatever it is that's wrong with the Moorean sentence."[31] We should thus be careful when encountering an odd-sounding proposition that has the form "p, but I don't ϕ p" and haphazardly calling it Moorean.[32] Consequently, I will only consider forms such as (1)–(7) as uncontroversially Moorean.

---

[31]DeRose (1991, p. 597); also in Douven (2006, p. 476).
[32]For a related discussion, see de Almeida (2007).

6.5 Moore's Paradox and Norms of Assertion

Now that we have a better understanding of the pathology of Moorean propositions, we are better positioned to discuss how the various proposed norms of assertion attempt to explain what makes Moorean propositions unassertible. In this section I present the typical strategies proponents of a range of norms of assertion have given for how each candidate norm can adequately explain the Moorean data. I will argue that KNA and reasons-based norms such as Lackey's RTBNA and Douven's *Rational Credibility Rule* (RCR) are relatively easily able to explain the data. TNA is also able to account for the data. However, there may be a cost for TNA in potentially needing to invoke a potentially controversial concept of having reasons short of those that would produce knowledge given true belief. I will argue that each of the norms is ultimately able to adequately explain the Moorean data. Although many KNA proponents view Moore's Paradox as providing differential support for KNA, I will argue that it does not. Consequently, we should not consider Moore's Paradox as providing differential support for any norm; rather, we should treat it as an adequacy condition on any proposed norm of assertion. Therefore, in order for a proposed norm to be a plausible candidate, it must be able to well explain the Moorean data.

Although, as I will explain, any assertoric norm that considers the belief of the proposition asserted as a necessary condition is able to explain the ommissive version of Moore's paradox – p, but I don't believe that p – as well as the commissive version – p, but I believe that *not-p*. However, it has been argued repeatedly that only KNA can well explain the knowledge

164

version – p, but I don't know that p.[33] That is, reasons-based norms like RTBNA are at an explanatory disadvantage. But while I will admit that KNA can well explain the data, the question will be whether it is the only norm able to well explain the data. Here is the typical KNA story. If knowledge is the norm of assertion – one may assert p only if one knows p – then in order to properly assert a conjunction, one must know both conjuncts. So in asserting "p, but I don't know p" one must know both "p" and "I don't know p." We can represent the assertion as Kp & K¬Kp. In order to know that one doesn't know p, one must not know p. That is, K¬Kp → ¬Kp. But in order to assert the first conjunct, p, one must know p. So in order to properly assert "p, but I don't know p" one must both know and not know p, which is a contradiction.

The underlying idea might be made clearer through a formal derivation. Where 'Np' means 'p is asserted with normative propriety';

1. Np → Kp . . . . (KNA)

2. N(p & ¬Kp) . . . . Assumption: Moorean case (MP)

3. K(p & ¬Kp) . . . . (KNA, 2, MP)

4. Kp & K¬Kp . . . . (3, Distributivity)

5. K¬Kp . . . . (4, Conjunction Elimination)

6. ¬Kp . . . . (5, Knowledge Factivity)

7. Kp . . . . (4, Conjunction Elimination)

8. Kp & ¬Kp . . . . (6, 7, Conjunction Introduction)

9. ¬N(p & ¬Kp) . . . . (2-8, *Reductio ad absurdum*)

---

[33]See, for example, Unger (1975, pp. 258 - 259); Williamson (2000, p. 23); Adler (2002, p. 194); and DeRose (2002, pp. 180 - 181). It's worth noting that the norms of assertion literature tends only to focus on the ommissive knowledge version, and doesn't discuss the commissive version – p but I know that *not-p*. When I speak of the knowledge version of MP, I mean the ommissive form.

In short: KNA diagnoses Moorean cases as asserted in violation of the norm. As Weiner puts it, "if the second half of the sentence is true, then the speaker will not satisfy the norm of assertion for the first half."[34]

KNA is also well suited to explain the problem with the belief versions of Moore's Paradox. Once again, if knowledge is the norm of assertion, in order to assert a conjunction, one must know both conjuncts. So in order to assert "p, but I don't believe that p," (the ommissive form) one must both know that p and that one doesn't believe p. That is, Kp & K¬Bp. Since knowledge requires belief, the first conjunct implies that one believes p: Kp $\rightarrow$ Bp. And in order to know that one does not believe that p, one must not believe that p: K¬Bp $\rightarrow$ ¬Bp. So in order to properly assert "p, but I don't know that p" one must both believe and not believe p, which is a contradiction.

For the commissive form – "p but I believe that *not-p*" – the KNA story could be the following.[35] de Almeida (2011) offers the following as an argument for why commissive forms of MP cannot be known. Suppose that premises (i) and (ii) are justified beliefs for me.

 (i) Everything my doctor says about me is true.
 (ii) My doctor says that both, [P] I'm terminally ill, but I believe that I am not [I believe that not-P].
 (iii) Therefore, I'm terminally ill, but I believe I am not.

de Almeida argues that (iii) cannot be a justified belief. Here's why. Let's assume, for reductio, that (iii) is true and justified for me. Since belief distributes, from (ii), I believe

---

[34]Weiner (2005, p. 230).
[35]The central norms of assertion literature tends only to discuss the ommissive form.

that P, and I'm justified in believing that P. And for the same reason, from (iii) I believe that not-P, and I'm justified in believing that not-P. However, suppose we adopt the following principle of anti-incoherence:

Anti-incoherence: If you believe that P and you believe that not-P (at the same time), you cannot know either, nor can you be justified in believing either.[36]

From anti-coherence, if I'm justified in believing that P and I'm justified in believing that not-P, then I'm not actually justified in believing either. So, by reductio, I'm both justified in believing P and not justified in believing P, which is a contradiction. Since I lack justification for the commissive form of MP in (iii), I cannot know (iii). Therefore, by KNA, I'm not warranted in asserting (iii).[37] More importantly, no assertions expressing the commissive form of MP will be warrantedly assertible on KNA. So one potential consequence of the KNA analyses is that Moorean propositions, both ommissive and commissive, may never be, properly speaking, assertible.

KNA's capacity to easily and comprehensively explain both principal versions of Moore's Paradox is frequently taken as evidence supporting it – presumably over at least some of its rivals. DeRose writes, "one of the most important recommendations of the account [is] that it provides a nice handling of the knowledge version of Moore's [P]aradox and other troubling conjunctions.[...]And the knowledge account can also be used to explain other troubling conjunctions that I can't see how to handle without the account."[38] Compare this to Williamson, who writes, "the hypothesis that not only knowledge warrants assertion makes

---

[36]For a discussion and justification for this principle, see de Almeida (2007).
[37]This proof works even assuming that the premises are *known* rather than merely justified beliefs.
[38]DeRose (1991, pp. 180 - 181).

it hard to understand what is wrong with an assertion of [(3)]."[39] However, there have been at least two attempts at an explanation for how reasons-based norms can equally well explain the data.

Igor Douven has argued for what he calls the Rational Credibility Rule of assertion (RCR):[40]

RCR: One should assert only what is rationally credible to one.

The details of what he means by "rational credibility" are not important, but an approximation is that a proposition is rationally credible if it is sufficiently highly probable and is closed under known entailment.[41] He provides a two-pronged attempt at explaining the unassertibility of Moorean propositions with RCR as the norm of assertion. First, the odd-soundingness of Moorean propositions may merely be a function of how rarely we encounter them. That is, it may not necessarily indicate an inconsistency even in terms of rational commitments. Recall DeRose's claim at the end of Section 6.4: we don't have a special faculty for distinguishing between sensing an inconsistency and whatever is wrong with Moorean assertions. However, this isn't a very convincing argument. After all, there is something persistently pathological with Moorean propositions that doesn't abate with more familiarity: they continue to sound odd even to those who work on them every day! For example, I've never heard someone assert, "I made pumpkin pie from purple pumpkins," but it doesn't sound distinctly odd the way that Moorean assertions do. So while it is an important reminder that we should be

---

[39]Williamson (2000, p. 253).

[40]Douven (2006). He also suggests a rule along the lines of "One should assert only what one can rationally believe."

[41]He includes a third condition that avoids the lottery paradox.

cautious when confronted with odd-sounding data, since the odd-soundingness may merely be a function of unfamiliarity, Douven's first prong doesn't appear to be a convincing analysis of the apparent unassertibility of Moorean sentences.

Douven's second prong is to explain directly the unassertibility of Moorean assertions using a reasons-based norm of assertion like RCR. If we suppose that rational credibility is the norm, then in order to assert a conjunction, one must have rational credibility for each conjunct. So in order to assert an instance of "p, but I don't know that p" one must have rational credibility for both p and that one doesn't know that p. It will be rare, if not impossible, for this to happen. Under what circumstances will it be rationally credible, by the speaker's own lights, that p is true and that she doesn't know it?[42] It's difficult to conjure up an example, but Douven suggests that an agent who can rationally credibly believe that she is in a Gettier situation may have rational credibility to believe p and that she doesn't know p because of her rationally credible belief that she is in a Gettier case (assuming that the agent knows about Gettier conditions, of course!). But consider also cases of psychoanalysis, in which one feels enfranchised to assert that p on one's therapist's say-so, while recognizing that one doesn't (fully, unambiguously) believe it oneself. For example, Jim strongly believes that his colleagues are conspiring against his tenure bid, but he also feels warranted in asserting that his colleagues are *not* out to get him on his therapist's say-so, even though Jim doesn't fully believe it. Alternatively, consider cases of temporary lucidity during a drug-induced hallucination, in which one is intuitively warranted in asserting that p then and there, but

---

[42]Or, for the commissive form, that it could be rationally credible, by the speaker's own lights, that p is true and that she knows that *not-p*?

might take one's (known) cognitive instability to preclude the satisfaction of subjunctive epistemic safety or sensitivity conditions on knowledge. However, one might have to reject the principle of anti-incoherence invoked above, thus identifying a weakness with this sort of example.

Jennifer Lackey, who argues for the Reasonable to Believe Norm of Assertion (RTBNA), also argues that reasons-based norms can well explain the Moorean data.[43]

RTBNA

1) One [may] assert that p only if it is reasonable for one to believe that p, and

2) If one asserted that p, one would assert that p at least in part because it is reasonable for one to believe that p.[44]

Like Douven, Lackey argues that a Moorean proposition satisfying RTBNA will be rare; however, she finds it easier to generate examples. A central class of examples in her case for RTBNA are what she calls selfless assertions. Her best example, I think, is her "Creationist Teacher" case.[45]

CREATIONIST TEACHER: Stella is a devoutly Christian fourth-grade teacher, and her religious beliefs are grounded in a deep faith that she has had since she was a very young child. Part of this faith includes a belief in the truth of creationism and, accordingly, a belief in the falsity of evolutionary theory. Despite

---

[43]Lackey (2007). Lackey's phrasing for 1) is "One should assert that p...." In keeping with the formatting of other norms such as KNA and CNA, I have changed it to "may."

[44]It's important to note that this is Lackey's formulation of the norm.

[45]Lackey (2007, p. 599).

this, Stella fully recognizes that there is an overwhelming amount of scientific evidence against both of these beliefs. Indeed, she readily admits that she is not basing her own commitment to creationism on evidence at all but, rather, on the personal faith that she has in an all-powerful Creator. Because of this, Stella does not think that religion is something that she should impose on those around her, and this is especially true with respect to her fourth-grade students. Instead, she regards her duty as a teacher to include presenting material that is best supported by the available evidence, which clearly includes the truth of evolutionary theory. As a result, while presenting her biology lesson today, Stella asserts to her students, "Modern day *Homo sapiens* evolved from *Homo erectus*," though she herself neither believes nor knows this proposition.

We can imagine a student of Stella's asking her whether *Homo sapiens* evolved from *Homo erectus*. Lackey argues that Stella could reasonably respond with either, "They did, but I don't know that they did," or, "They did, but I don't believe that they did." Intuitions may differ concerning whether Stella's assertion is too ridiculous to count as warranted. However, one could draw a distinction between warranted assertions (that aren't paradoxical or absurd in the Moorean sense), and warranted assertions that are still paradoxical or absurd in the Moorean sense. One could then say that Stella's assertion is warranted but still strikes us absurd. Now, if we assume that she really is making an assertion, and a warranted assertion at that, then selfless assertions of this kind would be a broad class of warrantedly assertible Moorean propositions.[46] But this is still consistent with the claim that assertible Moorean

---

[46]In Chapter 8 I argue that Stella is indeed making a warranted *assertion*.

171

propositions are rare.

Douven and Lackey thus suggest that there's no conclusive reason to think that it would be impossible to warrantedly assert a Moorean proposition. Reasons-based norms thus often admit that in some rare circumstances it may be permissible to assert a Moorean proposition. This is clearly a point of departure between KNA advocates and reasons-based norm advocates: KNA advocates tend to argue that Moorean propositions are never warrantedly assertible.[47] It is difficult to adjudicate, but in the absence of positive reasons to adopt the view that Moorean propositions are never assertible, I would err on the side of the more conservative claim that it may, in very special circumstances, be permissible to assert a Moorean proposition.

The truth norm (TNA) has a more difficult road in explaining the oddity and impropriety of asserting Moorean propositions. Since a defining feature of Moorean propositions is that they are consistent – that is, they can be true – if one may assert p only if p is true, then it seems that TNA licences Moorean propositions as properly assertible. But this is not what TNA advocates argue. Matthew Weiner, for example, has argued that the truth norm can adequately explain the Moorean data by arguing that TNA is not really the only norm operating on assertions. There are two aspects of this position. First, assertions are governed by Gricean norms such as the Maxim of Quality: *Do not say that for which you lack adequate evidence*; and, *do not say what you believe to be false.* So Weiner's truth norm is really a "reasonable true belief" norm. So what distinguishes this from KNA when most

---

[47]For example, Williamson (2000) and DeRose (2002). However, it's certainly possible for KNA advocates to argue that Moorean assertions are *typically* not warrantedly assertible, though not always.

epistemologists still largely agree that knowledge is (ungettiered) justified true belief? Weiner argues that "having adequate reasons" may still fall short of what is required for knowledge.

> For our assertions to be proper, not only must they be true, we must have reason to believe them true.[. . .] In many cases, the most likely warrant combined with truth will be enough for knowledge. When it is obvious that the speaker's warrant would not be enough for knowledge, assertion without knowledge is permissible.[48]

The second aspect of Weiner's approach is to utilize DeRose's primary/secondary propriety distinction.[49] Recall that for DeRose, an assertion is primarily appropriate if it satisfies the norm of assertion. An assertion is secondarily appropriate if the agent has good reason to think that their assertion satisfies the norm of assertion, even if it doesn't. So assuming that Weiner's version of TNA is the norm of assertion, a true Moorean proposition may satisfy the norm in a primary sense, but a speaker may fail to adequately satisfy the norm in a secondary sense. It may be true that both p is true and that the speaker doesn't know p, but they may fail to have adequate reasons to think that they're properly satisfying the norm of assertion. That is, the agent probably won't have good reason to think that their assertion is true even though it is true. Consequently, the speaker does not fully satisfy the norm. This will generally be the case for Moorean propositions: since they're potentially consistent, they will sometimes be true, but it's rare (if not impossible) for a speaker to have adequate reasons to think that they're satisfying the TNA norm in a secondary sense. However, like Douven and Lackey, Weiner might be open to the position that Moorean propositions are sometimes,

---

[48]Weiner (2005, p. 238).

[49]DeRose (2002) and elsewhere. C.f. Lackey (2007) for her rejection of this distinction. I also discuss the matter in more detail in Chapters 7 and especially 8.

though rarely, assertible, provided that a speaker could satisfy TNA in a secondary sense. TNA is thus able to adequately explain the Moorean data, including both the belief and knowledge versions.

6.6 The Moorean Moral for Norms of Assertion

I have defended the view that each of the principal candidate norms of assertion can adequately account for the Moorean data. This neutrality is not entirely neutral in its effects, though, since we saw earlier that KNA advocates have tended to claim that their proposed norm has a distinctive advantage in explaining the logico-linguistic data arising from Moore's Paradox. This claim has been undermined, I believe. It is undeniable that KNA has the resources to explain the Moorean data. But so too does every candidate norm currently defended in the literature, I have argued.

In particular, KNA may be "unnecessary for explaining why Moorean absurd thoughts [and assertions] strike us as contradictory."[50] If we take knowledge to generally be ungettiered justified true belief, then one can fail to know for many different reasons. As Littlejohn aptly explains,

> If you take yourself not to know that p because you take yourself not to believe p, say, then we have a situation where your second-order belief is transparently falsified by a fact about your own mind. If you take yourself not to know that p because you take it to be that p is false, we can explain the incoherence by

---
[50]Littlejohn (2010, p. 22).

appeal to the truth norm. If you take yourself not to know that p because you have insufficient evidence by your own lights, we can explain this in terms of the evidential norm. [...] And, if you take yourself not to know because you take yourself to be in a Gettier case...the evidential norm can explain the incoherence of that attitude. It seems that we have our bases covered.[51]

The moral is that there are many possible explanations for just the "p, but I don't believe that p" data. And what I have argued above suggests that there are many possible explanations which each coherently explain all of the Moorean data, including the "p, but I don't know that p" data.

That each candidate norm has an adequate explanation of the Moorean data is consistent with one norm's being the best explanation once we factor in broader considerations. However, the Moorean data itself doesn't differentially support one norm rather than another. Moore's Paradox taken on its own simply won't adjudicate between proposed norms; we need to appeal to a candidate norm's ability to explain other data, such as the Lottery Paradox (see Chapter 5) and prompts and challenges (see Chapter 7). We will need to appeal to a norm's ability to coherently explain a wide range of data, which will be discussed subsequently throughout this dissertation. Thus we should view a norm's explaining the Moorean data merely as an adequacy condition, not a deciding factor.

By eliminating Moore's Paradox as a source of differentiating evidence, this result heightens the significance of other data discussed in the norms of assertion literature. However, the need to accommodate Moorean cases may still function as a constraint on how candi-

---

[51]Littlejohn (2010, p. 97).

date norms can explain other data. As Matthew Benton suggests, "any explanation of what is problematic about the Moorean conjunction[s] ought also to explain why the challenge questions [such as "How do you know?'] are so apt, and vice versa: and even better, they ought to be given the same explanation."[52] So the project is to seek an elegant and unified explanation of all the data. In the next chapter I turn to the discussion of the linguistic data of prompts and challenges to assertion such as "How do you know?" Here too KNA advocates have held that their view is uniquely suited to explain the data. I will argue that reasons-based accounts are equally able to explain these cases as well.

---

[52]Benton (2011, p. 686).

Chapter 7

Challenges to Assertion: Why "How Do You Know?" Doesn't Challenge a Speaker's Knowledge

[1]

7.1 Introduction

There have been three broad categories of linguistic data marshalled as evidence for and against various proposed norms for assertion: statements about lotteries such as "Your ticket will lose"; Moorean statements of the form "p, but I don't know that p"; and challenges to assertions, such as "How do you know?" The first two categories have received much attention

---
[1]A version of this chapter is from McKinnon (2012).

in the literature. However, there has also been a consistent, if quieter, focus on the evidential role of challenges in the debate over norms of assertion.[2]

This chapter has two principal goals and two subsidiary goals. The first principal goal is to argue that the challenge "How do you know?" need not be interpreted as challenging a speaker's knowledge. Rather, the challenge is better interpreted as directly challenging a speaker's reasons. The second principal goal is to explain how a reasons-based norm of assertion, such as Jennifer Lackey's Reasonable to Believe Norm of Assertion (RTBNA) can equally well explain the range of linguistic data involving challenges to assertion as the Knowledge Norm of Assertion (KNA).[3]

> KNA: One should assert that p only if one knows that p.[4]
>
> RTBNA: One should assert that p only if it is reasonable for one to believe that p.[5]

In service of these principal goals, this chapter has two subsidiary goals. The first is to defend the view that the linguistic data of challenges to assertion should count as evidence in deciding between competing theories of assertoric norms. The second is to argue that we can evaluate the content and normative significance of challenges based in part on what would constitute *wholly adequate responses* to challenges.

---

[2]For example, see Hawthorne (2004), Kvanvig (2009), Lackey (2007), Rescorla (2009), Stanley (2008), Stone (2007), Turri (2010c, 2011), Unger (1975), Williamson (2000), and Weiner (2005), among others.

[3]Of course, in Chapter 9 I will return to explain how my proposed reasons-based norm, SRNA, can explain the data.

[4]Williamson uses the stronger construction: "One must: assert that p only if one knows that p."

[5]Lackey also proposes a second condition: "If one asserted that p, one would assert that p at least in part because it is reasonable for one to believe that p."

7.2 Challenges to Assertions

We have many ways to challenge assertions, including:

> C1: You don't even believe that.
>
> C2: Why do you believe that?
>
> C3: What evidence do you have for that claim?
>
> C4: But that's false!
>
> C5: You don't know that.
>
> C6: How do you know?
>
> C7: Are you certain/sure?

It has been suggested that the propriety of such challenges may provide support for a particular norm of assertion. Williamson, for example, argues that challenges such as C6 have a default propriety in most circumstances, and that KNA best explains this fact.[6] The upshot seems to be that if a particular challenge is widely appropriate, then this is best explained by an assertoric norm that can explain the propriety of the challenge(s).

Jonathan Kvanvig raises a worry about this sort of argument.[7] He notes that other sorts of challenges, such as C1 and C7, are also intuitively appropriate in most circumstances. Thus, says Kvanvig, this sort of argument appears to prove too much. By parity of reasoning, the appropriateness of using C7 to challenge an assertion in many contexts seems to suggest that certainty is the norm of assertion (CNA).

---

[6]Williamson (2000).

[7]Kvanvig (2009).

CNA: One may assert that p only if one is certain that p.

But this may set the bar for assertion unreasonably high: we're rarely certain of what we assert. Conversely, such a strategy may prove too little. In many contexts it is appropriate to challenge an assertion with C1, so this seems to suggest that belief is the norm of assertion (BNA).

BNA: One may assert that p only if one believes that p.

But this seems to set the bar for assertion too low. Overall, then, Kvanvig's view is that the data on challenges do not uniquely support any particular norm. He infers that such data should simply drop out of the discussion, since they fail to discriminate between the main competitors. Below I will return to whether Kvanvig is right about dismissing the data altogether.

Kvanvig provides an important objection to a central argument in favour of KNA. It is often claimed that the general propriety of "How do you know?" is good evidence that knowledge is the norm of assertion.[8] So by reminding us that a wide range of challenges are also appropriate in most circumstances, the objection demands a response by the KNA advocate. Furthermore, anyone wanting to use the linguistic data of challenges to assertions as evidence for a particular norm of assertion will also need to respond. Fortunately, a response is available.

---

[8]Williamson (2000), DeRose (2002), Turri (2010c), and elsewhere.

180

7.3 Turri on Challenges and the Knowledge Norm

John Turri specifically takes up Kvanvig's concern about using the linguistic practice of challenging assertions as evidence for norms of assertion.[9] Turri argues that KNA can explain the range of data and, on balance, that it best explains the data. Assuming the evidential probity of abductive reasoning, or inference to the best explanation, this implies that the linguistic data of challenges can indeed count as evidence for assertoric norms. Turri begins by splitting challenges to assertions into three categories: those inquiring after knowledge directly (call these *knowledge challenges*), those inquiring after something weaker than knowledge (call these *weaker challenges*), and those inquiring after something stronger than knowledge (call these *stronger challenges*).[10] C5-6 are examples of knowledge challenges, C1 is an example of a weaker challenge, and C7 is an example of a stronger challenge. He then deploys a dual strategy: he argues that all weaker challenges also challenge knowledge (call this strategy S1); and he argues that stronger challenges are often inappropriate (call this strategy S2). I will discuss these strategies in turn.

One popular definition of knowledge depicts it as (ungettierized) justified true belief.[11] There are thus four distinct individually necessary and jointly sufficient conditions for Knowledge. Turri's first strategy (S1) is to note both that the challenges to knowledge can take the form of challenges to these necessary conditions for knowledge, and that Kvanvig's seemingly weaker challenges are of precisely this sort. C2 and C3, for example, challenge an agent's

---

[9]Turri (2010c).

[10]Although he does not give the name "knowledge challenges," I introduce it for clarity and completeness.

[11]Kvanvig (2009, p. 9) defines knowledge as "(ultimately) undefeated justified true belief." Or what Rysiew (2007) calls the "ho-hum" view of knowledge.

justification for her belief of p. Since justification is a necessary condition for knowledge, C2-C3 also challenges the agent's knowledge of p. In this way, Turri argues that the propriety of weaker challenges does not provide evidence specifically in favour of a weaker norm than KNA (such as RTBNA).

An important feature of S1 is that it generalizes. Suppose that one were to think that certainty is the norm of assertion (CNA). Any weaker challenge than certainty (*viz.*, belief, justification, knowledge, *etc.*) will challenge a necessary condition of certainty and, thereby, challenge certainty. The strategy also generalizes to allow weaker norms, such as RTBNA, to explain the propriety of challenges weaker than reasonable belief; namely, belief and reasons. RTBNA can thus use S1 to explain the propriety of challenges such as C1 and C2. I find this strategy convincing. It provides a powerful strategy for explaining the propriety of prompts and challenges weaker than a proposed norm of assertion. Consequently, I will utilize this strategy for how a reasons-based norm can explain weaker challenges than "reasonable to believe" required by RTBNA.

However, while not specifically addressed by Turri, the ability of a norm to explain weaker challenges using S1 depends on the range of appropriate weaker challenges only involving necessary conditions of the norm. If, for example, there are some weaker challenges than knowledge, but which challenge something that isn't a necessary condition of knowledge, then the KNA advocate cannot deploy Turri's S1 to explain the propriety of such challenges. Fortunately, though, it seems plausible that the range of weaker challenges only involves necessary conditions of knowledge.

Now, *if* we assume that knowledge is the norm of assertion, then S1 nicely explains the propriety of both direct knowledge challenges, such as C6, and weaker challenges such as C1. Turri's second strategy (S2) is to now explain how KNA can explain the propriety of stronger challenges and how the data is best explained, on balance, by KNA rather than CNA. The argument for this partly involves forestalling the objection from the CNA advocate that S1 actually supports CNA rather than KNA.[12] He begins by arguing that insofar as stronger challenges are appropriate, they are consistent with knowledge as the norm of assertion. He adopts a Moorean view of assertion which involves interpreting warrantedly asserting *p* as the agent's representing herself as having the authority to assert *p*. The idea is that a stronger challenge bears on whether the agent is *accurately* representing herself as having this authority. Turri argues that this is distinct from actually having the authority: being sure that one has the authority is different from having the authority. From KNA, then, to assert that *p* is to represent oneself as knowing that *p*, while to ask whether someone is certain (*viz.*, to issue a stronger challenge) is simply to ask whether she is accurately representing herself as knowing. This is consistent with KNA as the norm of assertion since knowledge is what licences one to assert and asking if someone is certain is merely asking if she is accurate in her being licensed. The latter is sufficient but not necessary for warrantedly asserting.

Although stronger challenges, insofar as they're conversationally appropriate, are consistent with knowledge being the norm of assertion, Turri also argues that stronger challenges are often inappropriate. Although weaker challenges and knowledge challenges have a de-

---

[12]However, one could admit that S1 really does *best* support CNA while still arguing that the linguistic data, on balance, most favours KNA. Turri, personal communication.

fault propriety in most circumstances, the same is not true for stronger challenges such as C7. Here he presses the intuitive impropriety of stronger challenges and the view that knowledge is "more closely connected to assertion than is certainty."[13] Part of the argument for this latter claim is that "certainty does not figure in as full a range of appropriate challenges as knowledge does."[14] In normal conversational circumstances, if John asks Mary for directions, then it would generally be inappropriate for John to challenge her assertion with "Are you certain of that?" Often such challenges are only appropriate in high-stakes circumstances where accuracy in what is asserted is important.

Further support for S2 is found in the intuitive propriety of various *prompts* for assertion. A prompt is when one asks for an assertion, e.g. P1

P1: Do you know what time the meeting starts?

Turri also draws a distinction between stronger and weaker prompts. A stronger prompt asks for something stronger than knowledge, as in P2.

P2: Are you certain what time the meeting starts?

A weaker prompt asks for something weaker than knowledge – from the argument above, this will be a necessary condition of knowledge – as in P3.

P3: Do you have good reason to believe that p?

And, of course, there are prompts that directly ask for the speaker to assert something they know, such as P1.[15] The crucial point, Turri maintains, is that "*stronger prompts*

---

[13]Turri (2010c, p. 4). However, intuitions differ on the general propriety of stronger challenges. Some might think that "Are you sure?" is as widely appropriate as "How do you know?"

[14]Turri (2010c, p. 4).

[15]However, as I will argue, P1 may not actually be a prompt for *knowledge.*

are inappropriate."[16] But this isn't entirely accurate, he allows: rather, stronger prompts "typically become appropriate only *after* an assertion has been made."[17] Thus, the wider propriety of knowledge prompts, compared to stronger prompts, is taken as evidence that assertion is more closely associated with knowledge than certainty. Therefore the prompts and challenges data better supports KNA than it supports CNA.

Turri's overall argument that the linguistic data of prompts and challenges, on balance, best supports KNA amounts to an inference to the best explanation, in that KNA allegedly explains the aggregate linguistic data of prompts and challenges better than its competitors (i.e., BNA, RTBNA, CNA, *etc.*). S1 is taken to demonstrate that KNA can explain weaker challenges as well as challenges to knowledge. Furthermore, Turri argues that KNA explains stronger challenges because, first, such challenges are consistent with KNA and, second, assertion is more closely associated with knowledge than certainty.

Finally, we should consider the ability of a candidate norm to provide a unifying analysis of the wide range of linguistic data arising from prompts and challenges as evidence for that norm. So while Kvanvig is correct that the general propriety of a challenge such as "How do you know?" is not, on its own, evidence for KNA, KNA's ability to comprehensively explain the range of data, from C1 – C7, should be treated as evidence that knowledge is the norm of assertion. And so Kvanvig's contention that the data of prompts and challenges to assertions should drop out of the debate over norms of assertion is too hasty.

Admittedly, KNA can well explain the linguistic data of prompts and challenges to as-

---

[16]Turri (2010c, p. 5), emphasis in the original.
[17]Turri (2010c, p. 5), emphasis in the original.

sertion. This is not in dispute here. But does this explanatory capacity mean that Turri's inference to the best explanation performs as advertised? That is, does KNA *best* explain the data? Turri writes that it "remains to be seen whether [KNA's] weaker competitors have anything as plausible to offer."[18] In what follows, I take up this challenge by demonstrating that KNA is not, as Turri implies, "the only game in town." In particular, grounds to think that RTBNA can explain the data at least as well as KNA, are grounds to doubt that the data of prompts and challenges evidentially supports a knowledge-based norm over a reasons-based norm.

7.4   Wholly Adequate Responses

In this section I propose a principle for the interpretation of the normative significance and the content of challenges by considering what would constitute *wholly adequate responses.*[19] I will argue that wholly adequate responses to challenges such as "How do you know?" and "Are you certain?", which at first appear to challenge stronger epistemic concepts than "reasonable to believe," generally consist in providing one's reasons, which is well explained by a reasons-based norm of assertion. This will position us, in Section 7.5, to explain how a reasons-based norm such as Lackey's RTBNA can equally well explain the prompts and challenges data as KNA.

Consider the following exchange between colleagues. James asks Mary, "Do you know what time the meeting starts?", to which Mary responds, "Yes." In many contexts Mary will

---

[18]Turri (2010c, p. 4).
[19]The phrasing is borrowed from Stone (2007).

have given an acceptable response. Suppose that James is Mary's secretary and he's asking whether she knows the time of the meeting: if she answers "Yes," then he is satisfied with the response; if she answers "No," then he will tell her the time of the meeting. But now suppose that James is asking Mary because *he* wants to know the time of the meeting and he thinks that Mary knows the answer. Turri suggests that "Yes" would be an unacceptable (complete) response in this case. While it may be playful or funny, in a number of circumstances it would be an uncooperative inadequate response to the prompt.[20] We can say, following Jim Stone, that it would not be a *wholly adequate response*.[21]

Continuing with our example, suppose that Mary answers, "4 p.m." Challenges such as C6 being appropriate by default in most circumstances, suppose that James then follows up with, "How do you know?" I argue that wholly adequate responses to such challenges only require Mary to give her reasons for why she thinks that the meeting is at 4 p.m. Mary will not have to demonstrate *that* she knows; rather, she will merely provide her reasons, evidence, or justification.[22] She may respond, for example, "I read a memo and a poster indicating that the meeting starts today at 4 p.m." The epistemic term "know" (or its cognates) figures nowhere in the response, and the only epistemic concept apparently involved in the response is the speaker's reasons for her belief.[23]

Some writers have suggested that the challenge "How do you know?" imposes on responses

---

[20]Turri (2010c, p. 4). One could also invoke Gricean norms as an explanation. It's similar to answering "Yes" to "Can you pass the salt?" without also passing the salt. It's often funny, but annoying.

[21]Stone (2007, p. 97) proposes a principle for the interpretation of a question is to determine what a "wholly satisfactory answer would be."

[22]For the remainder of this chapter, and the dissertation, I will only refer to "reasons" but I will use it interchangeably for justification and evidence.

[23]The objection will be that the speaker ought not respond unless she actually knows the content of her assertion, but that is getting ahead of ourselves.

like Mary's the presupposition that the speaker actually knows.[24] The argument seems to be that by Mary straightforwardly answering the question, rather than demurring that she doesn't know, she has implicitly accepted the presupposition and means to communicate with her response that she knows. (I'll refer to this as the "presupposition view.") The point is more clearly observed had James asked Mary, "Are you absolutely certain, beyond any doubt, that the meeting starts at 4 p.m.?" Were Mary to straightforwardly respond, she would appear to communicate that she's certain and doesn't have any doubts about the time of the meeting. However, were she to have any doubts, she would demur and respond, "I think it's at 4 p.m., but I'm not absolutely certain." Advocates of the presupposition view thus argue that speakers' straightforward responses to prompts and challenges communicate that their responses meet the epistemic standards expressed in the prompt or challenge. That is, a straightforward response to "Do you know whether p?" will communicate that one claims to know that p, and a straightforward response to "Are you absolutely certain that p?" will communicate that one claims to be certain that p.

I think this is too strong a conclusion to draw from the evidence. We should first recognize that the evidence is being drawn from a number of everyday exchanges between speakers. But the folk use of "know" and its cognates isn't nearly as strong and specific as the epistemologist's concept depicting it as (ungettiered) justified true belief. One can walk into an undergraduate philosophy class and find students offering definitions of knowledge which don't require the truth of the proposition known. And elsewhere we can find evidence of definitions where one can know without good reasons, or even when there is overwhelming

---

[24]Williamson (2000, p. 252). Turri has also expressed this view in personal communication.

counter-evidence (through "blind" faith, for example). Moreover, the folk sense of "sure" and "certainty" isn't epistemic certainty. When we ask "Are you sure about that?", an affirmative response does *not* communicate that the speaker takes herself to be epistemically certain. Often it merely communicates that the speaker is very confident in the assertion.

Unfortunately, the presupposition view advocates may be importing highly theorized definitions of the apparently epistemic concepts being deployed in prompts, challenges, and responses. They may be assuming, for example, that the folk use of "know" is the epistemologist's definition of knowledge. We can find some evidence in how, often overly cautious or playful, epistemologists may plausibly have the same conversation that James and Mary are having. Epistemologist Mary, may respond to Epistemologist James' prompt by demurring that, well, she doesn't *know* what time the meeting starts because, after all, she may be in unknown Gettier conditions.[25] But often when regular speakers straightforwardly respond to "Do you know whether p?" prompts, they're merely expressing a confident belief, rather than the *epistemic* state of knowledge. Speakers in everyday contexts simply aren't sensitive to the highly specific epistemic sense of "know" and its cognates.

However, this is not to say that the presupposition view is without merit. There are certainly cases where a prompt or challenge creates a presupposition and the speaker's response is sensitive to this. In some cases if James were to ask Mary, "Can you tell me, for sure, what time the meeting starts?" she may plausibly demur and not straightforwardly respond because she doesn't want to communicate that she's sure of the meeting time. Moreover, we do notice that when speakers are made aware of the epistemic (or perhaps pragmatic)

---

[25]To the constant consternation of my colleagues.

implications of their assertions, they can be prompted to qualify the assertion. If I ask a stranger, "Do you know when the bus comes?" she may respond, "Yes, it comes at 5:12 p.m." Were I to challenge her assertion with "Okay, but do you *know* that it comes at 5:12 p.m.?" it's plausible to suppose that she may demur, "Well, maybe I don't *know*, but it says right here on the schedule that it comes at 5:12 p.m." But what explains the general willingness to straightforwardly respond to the "Do you know?" prompt, only to demur when challenged with focal stress on "know"? And what explains the observation that regular speakers tend to demur from stronger prompts, but tend to straightforwardly respond to knowledge prompts with assertions?

The answer is that "Do you know whether p?" doesn't generally create the presupposition, imposed on the speaker, that a straightforward response communicates a claim to know. It's simply not clear that "Do you know whether p?" is a prompt for the speaker to assert something she knows. Jim Stone observes that, "It's helpful to recognize that the question 'Do you know whether p?' needn't be a request for knowledge, or if it is, it isn't necessarily for the knowledge that p."[26] Such prompts often only solicit speakers to assert something they have good reason to think is true. The stranger has good reason to think that the bus comes at 5:12 p.m. and hence her response. When it's made clear that I'm soliciting a knowledge claim, she demurs. This shift is evidence against the presupposition view.

More evidence for my claim comes from properly analyzing challenges, such as "How do you know?", to assertions offered in response to "Do you know?" prompts. Let's suppose for the sake of argument that James' knowledge prompt does create the presupposition that a

---

[26]Stone (2007, p. 96).

straightforward response by Mary means to communicate that she knows. When James then challenges her assertion with "How do you know?" we would expect her to defend her claim to know, since she's communicated a claim to know that the meeting starts at 4 p.m. Since knowledge has four individually necessary and jointly sufficient conditions, we would expect her defence to include demonstrating that her claim satisfies each of these conditions. But this is not what we find.

Let's suppose that providing one's justification (what I've been calling "reasons") and evidence that one really believes the asserted proposition is simple. And putting the issue of demonstrating the truth of a proposition aside, suppose that in order to establish that one knows some proposition p, one must establish that one is not in a Gettier case.[27] But this is clearly not manifested in the linguistic data of adequate responses to challenges such as C6. Mary does not give reasons for why she isn't being Gettiered in her belief that the meeting starts at 4 p.m.; rather, she simply provides her reasons for believing that the meeting starts at 4 p.m. She tells James, for example, that she saw a memo and a poster, and has no reason to doubt their accuracy. In fact, it's plausible to suppose that she can be confident in her belief and yet lack evidence that she's not being Gettiered. Furthermore, the belief that one is not being Gettiered rarely, if ever, enters an agent's mind, let alone figures as part of a response to "How do you know?" Since the absence of Gettier conditions is the principal difference between a justified true belief and knowledge, if one were required to demonstrate that one knows in order to adequately respond to a challenge of "How do you know?" we should expect a speaker to provide evidence that they are not in such conditions. But, again,

---

[27]Lackey (2007) makes a similar point.

this is not what we find.

Consequently, I argue, wholly adequate responses to knowledge challenges such as "How do you know?" only require a speaker to offer her reasons for the assertion. Adequately responding to such challenges doesn't require the speaker to demonstrate that she knows. In further understanding why this is not the case, the "how" in "How do you know?" is critical. When we ask a speaker how they know, we're directly soliciting (epistemic) reasons for their claim. We're not asking them to demonstrate that they know, as evidenced above.

Since a wholly adequate response in many such cases is merely to provide one's (epistemic) *reasons*, the challenge expressed in C6 is best interpreted as a challenge regarding one's reasons. While C6 appears to challenge a speaker's knowledge, a wholly adequate response only requires the articulation of one's reasons. In Mary's case, she responds to James' challenge by citing her evidence for the meeting's starting time: she saw a poster and a memo.

One might object that since the question "How do you know?" features the term "know," it would be very odd if a proper answer didn't involve knowledge. In response, there are several points to make. First, as discussed above, the question involves the term "how," which seems to solicit reasons, rather than soliciting a knowledge claim per se. Second, compare such challenges with questions such as, "How do you throw a curveball?" Some people, such as pitching coaches, will be skilled at providing detailed answers that would put the questioner in a position to know how to throw a curveball. But there are other people, perhaps a naturally gifted athlete with a big league curve, who will simply shrug

their shoulders and say, "I don't know, I just can." We don't doubt that the player knows how to throw a curveball: look, he just threw one! We just doubt whether he can explain *how* he knows.

Malcolm Gladwell (2005), for example, discusses cases in which people reliably perform cognitively complex tasks, such as spotting forgeries, without always being able to articulate how they're able to do so. In some cases, such people can point to clear features of a forgery and say that such-and-such a feature is a tell-tale sign of a fake. But the term "know" and its cognates, and the epistemic concept of knowledge itself, figure nowhere in such responses. The responses merely consist in giving one's reasons. When we answer "How do you know?" with something along the lines of, "I know because [x]" we substitute our (epistemic) reasons for x. So we needn't be surprised that some responses to "How do you know?" questions do not involve the articulation of knowledge. Of course, the responses to such questions are importantly related to knowledge in that they require the speaker to articulate her reasons for the assertion, and these reasons are a necessary condition for knowledge, but the response need not itself involve knowledge claims.

So far I have argued that when faced with a challenge of "How do you know?" a speaker provides a wholly adequate response by merely providing her (epistemic) reasons. Examples abound.

James: "When does the meeting start?"

Mary: "4 p.m."

James: "How do you know?"

Mary: "I read a memo and saw a poster."

Steve: "Is the cafeteria serving pasta today?"

Linda: "Yes."

Steve: "How do you know?"

Linda: "I just walked by and saw it on the menu."

In a broad range of cases speakers adequately respond to such challenges simply by providing their reasons for believing what they assert. But I think this is true even for stronger challenges such as "Are you certain?" That is, as with knowledge challenges, wholly adequate responses to stronger challenges generally consist in merely providing one's reasons.

Now, one might think that stronger challenges require a stronger response than merely offering one's reasons for the assertion. But consider a case where a speaker is challenged with C7. Imagine a witness to a robbery and murder is being interviewed by the police. The witness has just provided a verbal deposition to the effect that she saw someone, whom she just picked out of a police line-up, rob a liquor store and shoot the store clerk, killing him. The interviewing officer then challenges the witness' testimony with C9.

C9: Are you certain that you saw the defendant rob the store and shoot the clerk?

A natural response for the witness is "Yes." But let's suppose that this is not adequate, under the circumstances, and the officer repeats C9 with an emphasis on "*certain.*" The natural tendency would then be for the witness to answer "Yes" but then provide some support for

her claim. This would include, for example, that she's confident that she can recognize the perpetrator.

So what, then, is the speaker of C9 actually soliciting? The officer is aware that the witness' testimony will be used in the pending trial of the suspect. And since he's aware that she is making a strong claim which will be central in seeking a conviction, she had better have very good reasons for her assertion which will stand up in court. So perhaps the challenge should be better interpreted as an indirect means of asking the witness, "How can you be certain, sufficient for court purposes, that you saw the defendant rob the liquor store and shoot the clerk?" Given the principle of interpretation suggested above, we should now look for what would constitute a wholly adequate response to this challenge. I argue that in most circumstances we would find adequate a response that laid out what we would consider (very) good reasons for the context. For example, the witness has very good eyesight, had a very good look at the defendant that day, and has very good memory. Furthermore, suppose that the witness was able to pick out the defendant with confidence from a police line-up with five similar looking individuals. Short of requiring epistemic certainty of the witness' testimony, these would generally constitute very good reasons even in high stakes court cases.[28] Since there is no counter-evidence to the witness' claim such as motivation to lie, evidence of bad memory/eyesight, or a reliable alibi for the defendant, the witness' testimony would be well-supported and unimpeached. Therefore, I submit, the response of (merely) providing reasons would be a wholly adequate response to a stronger challenge such

---

[28]We shouldn't interpret C9 as seeking epistemic certainty. If we take the skeptic seriously, then epistemic certainty about such events would be impossible. It would be silly, then, for an interviewing officer to ask for something that he knows is impossible (thereby violating Gricean norms).

as C7 and C9. What distinguishes wholly adequate responses for knowledge challenges from the responses required for stronger challenges is that the strength of reasons required for the latter will generally exceed that of the former: we demand stronger reasons for stronger challenges, and are satisfied with weaker reasons for weaker challenges.

Let's take stock. In this section I have argued that challenges such as "How do you know?" and "Are you certain?" are best interpreted as directly challenging a speaker's reasons, rather than a speaker's knowledge or certainty. I have done so by appealing to the principle of wholly adequate responses: in nearly every circumstance, what constitutes wholly adequate responses to such challenges only requires the speaker to provide her reasons for the assertion. Consequently, a reasons-based norm such as RTBNA is especially well positioned to explain this: challenges to assertions require a speaker to articulate her reasons, and having adequate reasons is what RTBNA holds as the central feature of the norm of assertion. But the KNA advocate may object that this is not yet a problem, since Turri's first strategy (S1) in dealing with weaker challenges is to show how they challenge a necessary condition of knowledge and, therefore, also challenge knowledge. This is not in dispute. However, what is in dispute is that KNA best explains the linguistic data. In fact, Turri essentially offers an "only game in town" argument: KNA is the only plausible candidate norm of assertion able to comprehensively explain the range of prompts and challenges data. In what follows I will explain how a reasons-based norm, such as Lackey's RTBNA, is a viable competitor to KNA and can equally well explain the data.

7.5 Reasons-based Norms as Viable Competitors to the Knowledge Norm

In this section I explain how a reasons-based norm of assertion, such as Lackey's RTBNA, can explain the propriety of the same range of prompts and challenges as Turri argues only KNA is well suited to explain. I will not argue that RTBNA best explains the data; rather, I will simply present a reasons-based norm explanation of the prompts and challenges data. By singling out Lackey's RTBNA, I don't mean to suggest that it's the only reasons-based norm competitor to KNA. There are others such as the simple *Reasonable Belief Norm* (RBNA) – assert only what one reasonably believes – or Igor Douven's *Rational Credibility Rule* (RCR) – assert only what is rationally credible to one.[29] I choose RTBNA because it seems to be the strongest competitor.[30]

Since the norm requires one to assert what would be reasonable for one to believe, it can accommodate some cases where it's reasonable for a speaker to assert something she doesn't presently believe. Lackey calls these "selfless assertions." A doctor named Sebastian, for example, has excellent evidence that vaccines don't cause autism and asserts to his patients, "Vaccines do not cause autism." Normally he believes this, but he doesn't presently believe it. His child was recently diagnosed with autism in close proximity to receiving a vaccine. Consequently, Sebastian finds himself doubting, perhaps irrationally, the overwhelming evidence supporting the assertion. Insofar as we consider this a warranted assertion (I think

---

[29]Douven (2006). I also won't discuss whether other norms such as TNA, BNA, or CNA can adequately explain the data. My only purpose here is to offer a defence of reasons-based norms as competitors to KNA.

[30]My general strategy for dealing with the prompts and challenges data will work, *mutatis mutandis*, for any reasons-based norm.

we should) and insofar as the doctor can be plausibly supposed to not believe the assertion (a more controversial point), this is a case where a speaker warrantedly asserts something he doesn't believe. In such cases, belief isn't a necessary condition for warranted assertion. Consequently, RTBNA seems to have an advantage over RBNA by being able to accommodate such cases.[31]

In explaining how a reasons-based norm such as RTBNA can equally well explain the range of challenges to assertions, I begin with the easy part: explaining how a reasons-based norm can explain the propriety of weaker challenges such as C1-C3. Since I find Turri's first strategy (S1) convincing, I will deploy it here. RTBNA can explain the propriety of weaker prompts and challenges because to challenge a necessary condition of RTBNA is to challenge RTBNA itself. However, since RTBNA doesn't consider belief as a necessary condition, how can it adequately explain the general propriety of C1 – C2? Suppose that there are some rare cases such as Lackey's "selfless assertions" where belief is not required for warranted assertibility.[32] One could explain the general propriety of C1 – C2 by noticing that we typically expect speakers to believe what they assert, but it may not be a necessary condition. Since we typically, but not universally, expect speakers to believe what they assert, challenges such as C1 – C2 will be appropriate in most, but not all, circumstances, such as with selfless assertions.

---

[31]For similar reasons, I'd suggest, RTBNA has an advantage over Douven's RCR. In fact, I think that there are cases where one can warrantedly assert something that one believes isn't rationally credible. Suppose that a speaker asserts that "Obama will not win in 2012." Further suppose that the speaker thinks that there's some evidence for this claim, but not enough to be rationally credible. Consequently, she can't yet bring herself to believe it. Given the evidence, in such cases the assertion is reasonable for the speaker to believe (and so plausibly satisfies RTBNA) but wouldn't be rationally credible to the speaker (and so fails RCR).

[32]Lackey (2007). See 8.5 below, especially page 241.

In fact, one may even grant that such challenges are always appropriate without therefore concluding that belief is a necessary condition for warranted assertibility. That is, a challenge may be appropriate without being *decisive*. A challenge is decisive if it provides a *prima facie* argument for the challenged assertion being unwarranted. Suppose that Mary asserts to James that the meeting starts at 4 p.m., but she's just guessing and James knows this. Then his challenge "But you're just guessing!" would be decisive: it would provide a *prima facie* case for her assertion being unwarranted.

If we accept that selfless assertions are those for which belief isn't required for warranted assertibility, then suppose that Sebastian's patients challenge his assertion with, "But you don't actually believe that, do you?" Although the challenge seems appropriate in the context, it's not decisive since he could explain the evidence supporting the assertion and his extenuating circumstances for (irrationally) doubting its truth. Sebastian thinks that he and his patients ought to believe that vaccines do not cause autism, but he just can't bring himself to believe it at the moment due to the circumstances of his child's recent diagnosis. And so while the challenge is apt, the patients will consider the assertion warranted after hearing Sebastian's explanation. Consequently, I argue, RTBNA can adequately explain the weaker challenges of C1 – C3.

One central feature of reasons-based norms is that truth of assertions is not required for warranted assertibility. Although we typically expect warranted assertions to involve true statements, it's not a necessary condition. So how can a norm such as RTBNA explain the general propriety of truth-based challenges such as C4? We should first notice, I think, that

challenges such as C4 are not *prima facie* decisive. Suppose that Fred overhears Mary's assertion to James that the meeting starts at 4 p.m. Fred is responsible for setting the meeting times and he had recently decided to change the meeting time to 3 p.m., but no notice has gone out yet regarding the time change. Were Fred to walk up and challenge Mary's assertion with, "But that's not true," this would probably not be a decisive challenge to Mary's warrant for her assertion. A reasonable reply on her part need not involve retracting her assertion and conceding that Fred is correct: more probable is that she would request reasons from Fred. "Oh really? Why is that?" she might say. One rarely takes a bare challenge as sufficient reason to reject an assertion.[33]

KNA and truth norm (TNA) advocates regularly argue that although Mary would be *blameless* for her assertion, she would still have violated a norm of assertion.[34] We wouldn't blame her for asserting that the meeting starts at 4 p.m. since her evidence is generally sufficient for justifiably believing and warrantedly asserting that the meeting starts at 4 p.m. However, she has somehow failed at achieving the goal of assertion in asserting something false. Keith DeRose makes a distinction between one actually satisfying the norm of assertion and one having good reason to think that one is satisfying the norm of assertion.[35] He calls the former *primary propriety*, and the latter *secondary propriety*. So, the KNA advocate would argue, a reasons-based norm can't explain the general propriety of challenges based on falsity such as C4. Like Lackey, I am highly skeptical that there's any meaningful distinction

---

[33]Though, if she knows that Fred is the person in charge of the meeting time, she might take his challenges as *prima facie* decisive. However, consider cases where the person issuing the challenge isn't in such a position of authority. If a professor asserts, "The heaviest element is X," and is an expert in the field of nuclear physics, then a student's challenge, "But that's not true," wouldn't be *prima facie* decisive.

[34]DeRose (2002) would be the most prominent example.

[35]DeRose (2002).

200

to be made between primary and secondary propriety. She writes, "the notions of secondary propriety and impropriety are spurious: either a speaker is behaving appropriately and is not subject to criticism *qua* asserter, in which case she has not violated a norm of assertion, or she is behaving inappropriately and is subject to criticism *qua* asserter, in which case she has violated a norm of assertion. There is simply no room for acts being secondarily proper or improper."[36]

We can ask, based on her evidence, should Mary have asserted to James that the meeting starts at 4 p.m.? I think that we should answer in the affirmative. Her evidence is good and she has no reason to doubt its accuracy. In fact, we may further suppose that she had recently confirmed with Fred that the meeting was still scheduled for 4 p.m. If, in the intervening time, Fred changed his mind unbeknownst to Mary, there is no meaningful sense in which we can criticize her assertion *qua*asserter. Sure, in a very weak sense we can "criticize" the assertion for being false, because *ceteris paribus* true assertions are better than false ones, but given her evidence, we think that Mary asserted properly. While it is certainly better for an agent to have good reasons for their assertion and for the assertion to be true than it is for the assertion to be reasonable but false, it doesn't follow that truth is a necessary condition for warranted assertibility. The upshot is that when a speaker has adequate reasons for her assertion, even if the assertion is false, challenges such as C4 are not decisive: having good reasons is generally an adequate defence from the challenge. That is, giving one's reasons is also a wholly adequate response to such challenges.[37]

---

[36]Lackey (2007, p. 604).

[37]There's a further question of what a speaker should do if it's been shown that her assertion was false, even if the speaker had good reasons. One might argue that she should retract or even apologize for her assertion.

While false assertions where speakers have excellent evidence for the truth of an assertion are blameless, and we may even think that such speakers assert properly, when speakers have no evidence and assert falsely or, worse yet, lie, such speakers are blameworthy: they've broken a rule of assertion. But what is the rule? That is, are the assertions poor because they're false? I think that this would be a misdiagnosis. The assertions are blameworthy because the speakers lack good (epistemic) reasons for what's asserted. But this is not yet to explain the propriety of challenges such as C4. Just as with the explanation above for the general propriety of belief-based challenges, in that we generally expect speakers to believe what they assert, we generally expect assertions to be true. Moreover, we expect speakers to attempt to assert truthfully. This is partly because what it means to have good (epistemic) reasons for a proposition is that an agent's belief properly aims at truth. Justified beliefs tend to be true because that's what it means to be justified. So a reasons-based norm can explain the general, or even universal, propriety of challenges such as C4[38].

7.6 Conclusion

In this chapter I've offered arguments in the service of two principal goals. First, "How do you know?" doesn't necessarily challenge a speaker's knowledge. Instead, such challenges can be plausibly interpreted as directly challenging a speaker's reasons. Second, a reasons-based norm such as Lackey's RTBNA can explain the range of linguistic data arising from prompts

---

One could alternatively argue that the speaker need not retract or apologize, but rather simply cease making the assertion. These are interesting questions, but would take us beyond the scope of the present project.

[38]Furthermore, a reasons-based norm is well positioned to explain why some such challenges are decisive while others are not.

and challenges to assertions at least as well as KNA. In Section 7.5 I offered arguments for how a reasons-based norm can explain the propriety of weaker challenges ranging from C1 to C4. I've done so by appealing to the principle of wholly adequate responses. In Section 7.4 I discussed how giving one's reasons for an assertion constitutes wholly adequate responses to knowledge and stronger challenges. A reasons-based norm is well suited to explain this since such norms require a speaker to have adequate reasons for assertions. The quality of the reasons required will change depending, in part, on whether a weaker, knowledge, or stronger challenge has been issued. Weaker challenges require weaker reasons, and stronger challenges require stronger reasons. So a reasons-based norm can explain the complete range of challenges from "Do you believe that?" to "Are you certain?" Consequently, a reasons-based norm can comprehensively explain the same range of prompts and challenges data as KNA. Furthermore, I think that a reasons-based norm can explain the data at least as well as KNA. A unified explanation is available based on the principle of wholly adequate responses. Turri writes that it "remains to be seen whether [KNA's] weaker competitors have anything as plausible to offer" as an explanation for the linguistic data of prompts and challenges.[39] I submit that it no longer remains to be seen: reasons-based norms, such as Lackey's RTBNA are viable competitors and KNA is not the only game in town.

---

[39]Turri (2010c, p. 4).

Chapter 8

Betting on Assertion: The Supportive Reasons Norm of Assertion

8.1   Introduction

In this chapter I present my proposal for the central norm governing the practice of assertion, which I call the Supportive Reasons Norm of Assertion (SRNA). I argue that by properly understanding the relationship between the goal of assertion and its nature as a social practice, we can recognize that there are conventional and pragmatic features *internal* to the practice. These internal features should be made explicit in our articulation of the norm, and not treated as inessential to it. This feature of my view sets it apart from other extant proposals of reasons-based norms. In making my case I suggest that a useful analogy can be drawn with properly understanding the relationship between the goal and norm of betting (*i.e.*, placing wagers). Insofar as the results are irrelevant to the normative assessment of a

bet, I argue that the same applies to assertion *viz.* truth (or knowledge). I therefore argue against factive assertoric norms such as TNA and KNA.

I then discuss a recent view of the norms of assertion offered by Ernest Sosa who, like my own view, argues that we should understand normative assessments of assertions in terms of performance normativity, and that the norms of a practice derive their content from the goals. However, he takes these considerations as supporting KNA, rather than a reasons-based norm. Since my view shares these central features, but I argue for a reasons-based norm, Sosa's view provides a useful foil. I'll argue that taking such a view about the nature of norms and assertion best supports a reasons-based norm rather than KNA.

The chapter will begin with a brief presentation of SRNA. Section 8.3 presents the analogy with norms of betting, with an aim of defending a norm of assertion where truth (*i.e.*, "success") is not a necessary condition. I then turn to more clearly explicate a number of key concepts of SRNA not addressed in Section 8.2. Section 8.5 presents Sosa's view, including his idea of meta-aptness, as a foil for my view. I use it to draw out important distinctions in SRNA, which clearly sets it apart from other reasons-based norms.

8.2 The Supportive Reasons Norm of Assertion (SRNA)

I propose the following norm as the central constitutive norm for the linguistic practice of assertion:

SRNA:

(i) One may assert that p only if the speaker has supportive reasons for p, and,

(ii) The relevant conventional and pragmatic elements of the context of assertion
are present.

Borrowing some analogous terminology from Kaplan, we can say that SRNA has an invariantist normative *character* but a contextualist normative *content*.[1] These approaches are familiar from existing literature, but they tend not to go together. I unify them by making the invariant component a type of indexical; the normative character is a function from context to content. Simply put, the context of an assertion determines the precise character of the norm for that context. Conditions (i) and (ii) are individually necessary and jointly sufficient conditions for warranted assertibility, and this does not shift with context. Hence, the structure of the norm – the character – is invariant. There must be supportive reasons for the proposition asserted (I will say more on what this means below) and it is always the case that the relevant conventional and pragmatic elements must be present. However, what count as supportive reasons for a given proposition, what counts as relevant, and what count as conventional and pragmatic elements possessing that relevance, will each shift with context. The normative content of SRNA is thus highly contextual.

8.3 Betting as a Model for Norms of Assertion

I will put off explicating these many concepts and features of SRNA until Section 8.4. Before I do that, I propose that a useful analogy for understanding the goal and norm of assertion can be found in understanding the goal and norm of betting. The purpose of the analogy

---

[1]Kaplan (1989).

is to argue against factive norms of assertion – ones on which truth is a necessary condition for properly performing an assertion. A number of writers have argued for a non-factive norm of assertion, most notably Douven (2006), Lackey (2007), Hill & Schechter (2007), and Kvanvig (2009). However, these views are often advanced through the use of carefully selected examples. While that approach has merit – and to a certain extent, my own view depends on key examples – my argument against factive norms involves properly appreciating a critical distinction between the norm and goal of a practice, and, moreover, the connection between the goal of a practice and its norm.[2] The key point will be: one can place a good bet and lose, just like one can perform a warranted assertion that happens to express a false proposition.

There is a rich analytical literature on betting communities, particularly of professional poker players, involving how people ought to go about being professional gamblers.[3] In the last few decades there has been an explosion of interest from mathematicians, decision- and game-theorists, and psychologists, all contributing to the theoretical analysis of poker and its norms. Professional poker is no longer a place for intuitions and gut-feelings, but rather careful analysis of betting patterns, calculating probabilities, and game-theoretically optimal strategies of play. Moreover, there is a body of research focusing on how professional poker players should analyze their results, manage their bankrolls, and generally think about the practice of betting itself.

Critical in discussing norms of professional gambling is noticing systematic differences

---

[2]For the moment, I set aside the question whether there are multiple goals or norms of practices.

[3]While I discuss the norms for professional players, these norms are in force for anyone playing poker with the goal of winning money.

between how people actually tend to behave and how people ought to behave. For example, in the behavioural economics literature there is broad support for the observation that real people tend to be risk averse rather than risk neutral. Consider a toy example in which I offer you the following wager: Pay \$1 to bet on the flip of a coin. I will offer you \$1.02 if you guess correctly; you receive nothing for guessing incorrectly. We can calculate the expected value (EV) of this wager: the long term expected pay-off of such a wager. We do this by using the following formula.[4]

EV = (Probability of Winning * Amount Won) – (Probability of Losing * Amount Lost)

The EV of this particular wager is \$0.01 (a penny). In decision theory an agent is risk neutral if she would equally prefer to receive the EV of a wager to taking the wager itself. So, an agent would be risk neutral in this case if she would equally prefer to just accept a penny to taking the wager of betting the \$1 to win \$1.02.[5]

Being risk-neutral is a pillar of the practice of a successful professional gambler. And since risk-neutrality is one of the rational means for properly achieving the goal of poker (*i.e.*, making money), being risk-neutral is a norm of professional poker. There's a critical distinction to be made here between a norm and a rule of a game. It's a constitutive rule

---

[4]This is only for simple wagers with a single payoff.

[5]Of course, normal people aren't risk neutral. Instead, there is a general tendency to take the penny rather than risk losing the wager even though they are theoretically worth the same. Consider another example. Suppose that I were to offer you \$125 on correctly guessing the outcome of a coin toss to your \$100. The EV is \$12.50. It's generally observed that people would greatly prefer to take \$5 rather than accept the wager. That is, they'll accept something considerably less than the EV of the wager if it means avoiding the risk of losing. Although most people are risk-averse, some are risk-seeking. For example, they'd rather wager their \$100 to win \$125 rather than accept \$15 (possibly because they value making wagers for their own sake). However, the objectively rational position is to be risk-neutral, particularly if one is a professional gambler. See, for example, Hilger & Taylor (2007).

of Texas Hold'Em poker that the player to the left of the Big Blind starts play (pre-flop). However, this isn't a *norm*. Following Rescorla (2009), it's a *requirement* of playing poker that one follows the rules of the game. However, one need not follow norms to count as playing the game. In such cases, one merely plays it poorly. Likewise, one could fail to be risk-neutral and still count as playing *poker* by following the necessary requirements (i.e., rules). Hence, being risk-neutral is a norm, rather than a rule or requirement of the game.

Now, while the value of risk-neutrality is merely instrumental towards achieving the goal of making money, it would be a mistake to view the constitutive *norm* of professional poker as the winning of money. Naturally this is the principal *goal* of the practice, in that professionals gamble because they can make money; and, enjoyment aside, they gamble only because they can make money. While one might think that the real goal is winning pots (hands) or winning tournaments, this is a mistake. Both of these would be better understood as intermediate goals in service of the overarching goal of winning money. One aims to win hands so as to win money; likewise, one aims to win tournaments so as to win money (and possibly prestige). But, of course, people play poker for all sorts of reasons such as entertainment, to live a certain lifestyle, and so on. So why should we privilege winning money as the central goal of playing poker? By analogy, one generally thinks that the goal of playing soccer is to win, but one might play in a professional soccer game with a special ostentatious flair, with the intention of sending the political message that people of one's homeland will not be cowed by a repressive government.[6] That could be the goal of one's playing soccer. But the availability of the option to play soccer with that goal asymmetrically depends on the goal of playing to

---

[6]My thanks to Tim Kenyon for the example and phrasing.

win (or whatever the success conditional goal is); whereas the opposite is not true.

In poker, the goal of winning money asymmetrically defines the practice, and so we can say that winning money is the constitutive goal of poker, while recognizing that there are all sorts of parallel goals that people can pursue. Barry Greenstein, a famous poker pro and philanthropist, plays poker tournaments and donates his winnings to charity. He's now known as the Robin Hood of poker.[7] So he plays with the goal of donating to charity, but the availability of this option to play poker with the goal of donating money asymmetrically depends on the goal of playing to win money. However, the goal of playing to win money doesn't depend on the goal to donate to charity. This is why the goal of winning money has a special prominence in the practice of gambling (and poker, in our case), and I call it the constitutive goal of the practice. Of course, some practices might have multiple constitutive goals (just like a practice can have multiple constitutive norms),[8] but for our purposes I propose that there's a single constitutive goal of poker, and it's to win money.

It's important that we keep the norms and goals of a practice separate, though. For example, we shouldn't take the goal of winning money as the guiding *norm* of poker. To see why, imagine that we did take the norm of poker simply to be winning money. Then, were we to compare two players to determine who was best, we would only have to look at who had won the most money. But this produces a problem. Suppose that there are only two poker players at a table, Alan and Bonnie, and they only play one hand. Alan wagers all his money with one card left, and has only a 5% chance of winning. Bonnie has wagered all

---

[7]In fact, he was instrumental in beginning Team 1%, a group of players who pledge to donate 1% of their winnings to charity.

[8]Here we must remember the distinction between goals, norms, and requirements.

her money, and has a 95% chance of winning (assuming that the evidence is transparent to both players, after making the wager, and that no ties are possible). However, as luck would have it, Alan wins and now has all of the money. Is he the better player? Arguably, he has made a massive error (he got his money "in bad" as they say) and Bonnie has played very well. But the *results* appear to side with Alan. Since the role of luck cannot be eliminated from gambling, and luck is out of an agent's control, it is a mistake to focus on results as a decisive indicator of skill.[9]

In poker, mistakenly focusing on results that may be the consequence of luck rather than skill is described as being too *results-oriented*. If a gambler has a losing session, it would be incorrect to immediately infer that she made a mistake. It's quite possible for her to have made every decision correctly and still have lost, just as it's possible for her to have made every decision incorrectly and still have won. Consequently, well-informed gamblers focus on whether their choice of bets showed positive EV (and, ideally, that their choices *maximize* EV); simply because a bet was lost does not mean that it was a bad bet. Bonnie made the right bet regardless of the results.

In response to cases like Alan and Bonnie, one influential poker theorist has proposed a concept called "Sklansky Bucks" (named after himself, of course).[10] Sklansky Bucks (SBs) are a version of an EV calculation. They're the theoretical value of a play in poker regardless of the actual outcome (due to chance). So suppose that Alan and Bonnie have $1000 to bet and Bonnie gets her money in as a 95% favourite. We calculate her SBs by determining her

---

[9]I set aside discussions of luck, such as Pritchard (2005) and Lackey (2008) who argue against a "control" view of luck.

[10]Sklansky (2007).

equity (the probability of her hand winning: 95%) which means she theoretically "owns" 95%

of the total pot ($2000). We take this value and subtract her portion of the bet. Her bet is

therefore worth ((0.95 * $2000) - $1000) = $900. Her bet "earns" her $900 SBs regardless

of the results.[11] Even if she loses, we can say that her play earned her $900 SBs. Of course,

both EV and SBs can be positive, neutral, or negative. Alan's play "loses" him $900 (((0.05

* $2000) - $1000) = -$900). So although Alan actually wins the hand (and has $2000), we

can objectively rate Bonnie's play as much better than Alan's and attach values in terms of

SBs.

The importance of introducing the concept of SBs, in order to evaluate one's play, is

to avoid focus on actual results, since such results involve chance factors beyond a player's

control.[12] The moral being that normative evaluations should only take place on features

which the agent can control. Since the agent can control her play, but not the actual outcomes

of draws, we evaluate the former but not the latter.

The critical consequence of this is that norms of poker are at the level of evaluating SBs

rather than dollars won or lost. This means that while poker players are ultimately interested

in maximizing their winnings (the *goal* of the practice), they should act to maximize SBs

---

[11]This means that were she to win the $2000, since she only has equity for $1900, she would win "more" than she should to the tune of $100. Were Alan to win, since he only has equity for $100, he would win $1900 more than he should.

[12]One might think that SBs are just another sort of outcome or result. But this is mistaken. Take the poker hand example discussed above: in some cases, Bonnie can *know* that she has the best hand and that, at worst, she has a 95% probability of winning. For example, if Bonnie holds pocket Jacks to Alan's pocket Queens on a board with J952 (with no flush draws available to either player), then Bonnie can *know* that she has the best hand (at the moment) and the only way for a hand to beat her is for her opponent to have pocket Queens, Kings, or Aces, and to draw one of two cards (respectively) to beat her. Her decision to call is based on realizing this: she doesn't gain the "result" of earning $900 SBs in the same sense of winning or losing the actual hand. She performs the action of calling (*i.e.*, the decision) and immediately her play is evaluable in terms of SBs: we don't need to wait to see what card comes on the river.

(the *norm* of the practice).[13,14] Indeed, the norms of poker incorporate the maximization of SBs precisely because this is the overriding rational means of aiming at the ultimate goal of making money. A cautionary joke among poker players has it that you can't pay your rent with Sklansky Bucks; but there is also no rational way to attempt to pay the rent through poker without maximizing Sklansky Bucks. Hence I suggest that the constitutive norm of poker is to maximize SBs: this is what ought to guide the selection of moves in the game of poker.

Nevertheless, we could elicit an intuition about the norm of betting, similar to the intuition about the norm of assertion, that the norm should be factive. The idea is that betting in accordance with the governing norm consists in winning is analogous to the intuition that asserting in accordance with the governing norm is speaking truly (and perhaps expressing knowledge). However, I argue that this is a mistake in the betting case, and that the nature of the mistake helps to clarify why it is implausible in the assertoric case as well. Consider the following bet. I offer you $1,000,000 ($1M) to your $1 on the flip of a fair coin. It should be obvious that this is a terrific opportunity for you. The EV of the wager is nearly $500,000 with very little risk on your part (let us assume that it's not your last dollar, and that a single dollar does not buy very much of what matters to you). Suppose that you take the

---

[13]One might think that this claim is absurd since so few people who play poker satisfy the maximize SBs norm (indeed, only a subset of successful *professional* poker players live up to the norm). So, one might think, are all of these other people simply not playing the game? Of course not. As Williamson (2000) noted, one can assert many false things, thus failing to satisfy the knowledge norm of assertion, and still count as *asserting*. One can fail to satisfy the constitutive norm of a practice while still participating in the practice. So poker players who do not maximize SBs in their play are still playing poker, they're just not playing as well as they could.

[14]One might additionally think that maximizing SBs isn't a norm, it's a *strategy*. But this is an error: check-raising on a flush draw is a strategy (aimed at maximizing SBs), as is raising on a bluff 45% of the time in a situation (say, when you're on the button and an aggressive player to your right open raises). Maximizing SBs isn't a strategy: it's what it means to properly aim at the goal of winning money – it's the norm.

wager and lose; how should we evaluate your bet? Was it a good bet, or a bad bet? What is obvious is that you lost, so if the intuition is that only winning bets are in accordance with the norm of betting, then we must say that the bet did not live up to the norm. If there is a corresponding intuition to say that people made an error when making losing bets, then we must conclude that you ought not to have taken the bet. But this is clearly a mistake. We evaluate the bet not based on the actual results, since these are out of your control, but on whether your wager had a positive EV. When there is only the option to take a bet or not take a bet, and the agent is risk neutral, then she ought to accept the bet if and only if it has a positive EV. If there are multiple options between wagers (including not taking any wagers), then the agent ought to accept a bet if and only if it is the most positive EV option. In this case, since there is only the option to accept the bet or not, and the EV is (massively) positive, the agent ought to accept the bet. But this evaluation is completely independent from the actual results of the wager (based on the flip of the fair coin). Whether you win or lose is not up to you and so we do not evaluate it. Instead, we evaluate your decision to accept the wager (independent of the results) and whether it maximizes EV.

As discussed in Chapter 3, we can see how the norm of the practice of betting derives its content from being the standard of assessment for the rational pursuit of the goals of betting. "Wager that p only if the wager will win" isn't a plausible norm for being a professional gambler; placing winning wagers is, however, a plausible goal. The problem with equating winning (*i.e.*, being results-oriented) with positive normative assessments of agents is manifest in the above example of Alan and Bonnie. She did very well in going about attempting to

achieve the goal of the practice, and her falling short of achieving the goal doesn't clearly bear on the normative assessment of her actions, for the reasons mentioned.

And so I suggest an evaluative parallel between norms of assertion and norms of betting. We may conceive of knowledge or truth as the goal of assertion, but from this it does not follow that knowledge or truth is a requirement for proper assertion – *i.e.*, the norm.[15] In betting, the usual goal is to win money, but it is not a requirement of proper betting for one to make money. Instead, the constitutive norm of betting is maximizing SBs. The actual results are irrelevant. I argue that the same applies to assertion. Although there is a strong intuition that a successful assertion is one that expresses knowledge (and hence truth), this would be to mistakenly focus on the results of the assertion rather than whether it was a properly formed assertion. That is, I want to suggest that whether an agent satisfies the conditions for SRNA in making an assertion is sufficient for making a warranted assertion independent of whether the agent's assertion happens to express truth or knowledge. What we really care about in making proper assertions is having good reasons; having an assertion rise to the level of knowledge is merely a fortunate bonus, just like winning a positive EV bet.

So if we assume that the goal of assertion is truth, or even knowledge, how can asserting something that fails to achieve the goal still count as a normatively appropriate assertion? There are a number of cases, I think, where we consider false assertions as appropriate. In particular, there are some cases of warranted assertions that happen to be false due to

---

[15]I'll have more to say, in Section 8.5, on what plausible goals we can identify for assertion and what this means for plausible norms.

bad luck. Consider Bill who is a taxi driver very familiar with the local roads and popular destinations. After his shift, Bill's friend Penny asks him the quickest route to the concert hall. Bill has just driven the route an hour ago; traffic was good and there were no indications of looming problems, and he gives Penny directions accordingly. Normally we would take it that Bill knows the quickest route to get to the concert hall, and that what he's asserting to Penny is also an instance of knowledge. However, suppose that shortly after Bill's trip some road construction began that will greatly slow the traffic on the route he has recommended. That is, his assertion that the recommended route is the *quickest* route is false by the time he utters it to Penny.[16]

In this case, I argue, Bill is warranted in making his assertion despite its falsity and its failing to help Penny achieve her goal of arriving on time. He has every reason to consider his directions as correct, and doesn't have reason to think that a recent road closure is likely. And his being in this evidential situation is not the result of any failure of epistemic due diligence. In fact, we may even suppose that Bill has strong evidence that road conditions will not change after he recently took the route – say, his having seen a public works billboard describing the recent *completion* of construction on that stretch of road. Given his strong evidence for the accuracy of his directions, I suggest, Penny would have to agree that Bill asserted properly – that it would be unreasonable to have expected better of him – even though his directions didn't get her to the concert hall on time. Thus, I argue, false assertions

---

[16]Or consider an alternative example. Penny asks Bill for directions to the concert hall. Her question creates a presupposition on Bill's answer such that his directions communicate that he thinks that the concert hall is still where he thinks it is. Bill gives his directions, but unbeknownst to him, in the short time since he last drove by, it has been miraculously demolished. So his directions are to a concert hall that no longer exists.

may be warranted in cases where the speaker is strongly justified in believing the proposition asserted and the assertion is false merely due to bad luck.

A common response to examples of this sort, from the truth- and knowledge-norm advocates, is that although Bill's assertion isn't blameworthy, it still violated the norm of assertion because it was false. They argue, for example, that his assertion had *secondary propriety* (that is, he had good reason to think that his assertion satisfied the norm), but lacked primary propriety (because it violated the norm).[17] Some have questioned whether this is a meaningful distinction. For example, Lackey writes,

> [T]he notions of secondary propriety and impropriety are spurious: either a speaker is behaving appropriately and is not subject to criticism *qua* asserter, in which case she has not violated a norm of assertion, or she is behaving inappropriately and is subject to criticism *qua* asserter, in which case she has violated a norm of assertion. There is simply no room for acts being secondarily proper or improper.[18]

I think that Lackey's dismissal misses its mark. The distinction between primary and secondary propriety strikes me as important. We intuitively recognize many cases where agents fail to satisfy the norm of a practice, but we refrain from blaming them due to their reasonableness in thinking that their action satisfied the norm. Consider the case of Mary in Chapter 7. Suppose that there's a department meeting scheduled for 4:00 p.m. There was an e-mail memo circulated posting this time, along with a few posters announcing the meeting

---

[17]Compare Williamson (2000), DeRose (2002), Weiner (2005) and others.
[18]Lackey (2007, p. 604).

posted around the department. So suppose that James asks Mary what time the meeting starts, and she asserts, "The meeting is at 4:00 p.m." This should strike us as a warranted assertion. But suppose that, unbeknownst to Mary, the meeting time was recently changed to 3:30 p.m., but not announced. So Mary's assertion is false. If knowledge is the norm of assertion, then she asserted improperly, strictly speaking. Her assertion thus fails to be primarily appropriate. And *if* knowledge is the norm of assertion, this seems correct. Since we have the further intuition that Mary couldn't have known that the meeting was actually at 3:00 p.m., and she had very good evidence for her assertion, we don't blame her: her assertion was secondarily appropriate. So, assuming that knowledge is the norm of assertion, the primary/secondary propriety distinction has merit.

It seems to me that Lackey's point, and my point in Chapter 7, is that if we don't assume that knowledge is the norm of assertion, then there's a viable argument that Mary's assertion was fully appropriate. But it doesn't follow from this that the distinction between primary and secondary propriety is worthless. Suppose that RTBNA is the norm of assertion. And suppose that Mary* didn't read the e-mail memo and believes that the meeting is at 4:00 p.m. based on the last meeting being at 4:00 p.m. And while this is true, she's new to the department and doesn't know that nearly all meetings are at 3:30 p.m., and the last meeting being at 4:00 p.m. was a special case. And suppose that James is also a new member and thus finds Mary*'s reasons plausible. One might suppose that Mary* reasonably *thinks* that she has good evidence for her assertion, thus reasonably thinking that her assertion satisfies the RTBNA requirement for it to be reasonable for one to believe what she asserts. But from

Mary being reasonable in thinking that her assertion satisfies the norm, it doesn't follow that her assertion actually satisfies the norm. Given a broader understanding of Mary* being a new department member, her inference from a small sample size of department meeting times doesn't make it reasonable for her to believe that the meeting is at 4:00 p.m. The distinction is essentially the difference between subjective and objective justification: an agent is subjectively justified if, from her perspective, it's reasonable to believe (or assert) that p; an agent is objectively justified if it's *actually* reasonable for her to believe (or assert) that p. So even on Lackey's view of the norm of assertion, the distinction between primary and secondary propriety has merit.

However, it remains an open question whether Bill's assertion, just like Mary's, is warranted. To motivate this sort of view, I think we should ask whether Bill should have asserted what he did to Penny, given his state of information?[19] It's hard to make sense of our normal assertoric practices unless we answer "Yes." His evidence was strong, recent, and not impugned by available defeaters. Was his assertion perfect *tout court*? Obviously not, since it was false and led Penny astray. So while, *ceteris paribus*, a true assertion is better than a false assertion, and a known assertion is better yet, it doesn't follow that false or unknown assertions are unwarranted. It follows that epistemologists must distinguish between an assertion's being suboptimal, in that it doesn't rise to, say, knowledge or even certainty, and its being warranted. Examples of the sort I have canvassed show that an assertion can be warranted while still falling short of the best possible sort of assertion, or even short of many other "better" assertions.

---

[19]From an objective perspective analyzing the case, that is.

The upshot of my argument is that we can say that one assertion would be better than another without concluding that the worse assertion is unwarranted, or that the epistemic state attained by the better assertion is the norm of assertion. Let's return to the betting analogy. If Bonnie makes the $1000 bet with a 95% chance to win, then a bet that wins is, *ceteris paribus*, better than a bet that loses, but it doesn't follow that the losing bet is a normatively inappropriate bet. I argue that the case is the same for assertion. Analogously with the norm of betting, the norm of assertion should not be overly results-oriented. I'll have more to say about this in Section 8.5.

8.4   Explicating SRNA

I claimed in Section 8.2 that SRNA has an invariant character but a contextualist content. It's important to explicate what I mean by this. There's a debate in contemporary epistemology whether the truth conditions of statements involving "know" and its cognates to agents depends on context. Consider DeRose's famous bank cases.[20] The issue is whether an agent, call her Mary, can be said to know whether a bank is open on the subsequent Saturday in different pragmatic and conversational contexts. In the low stakes (LS) context, being wrong about the bank's being open is not especially important. But in the high stakes (HS) context, being wrong carries with it particularly bad consequences. In both contexts Mary has the same evidence for the proposition that the bank is open on the subsequent Saturday; namely, that she has seen it open on previous Saturdays and that banks tend to be open on

---

[20]DeRose (2002).

Saturdays.

The contextualist intuition is that Mary may be said to know in LS but not in HS. In the LS context, challenging her knowledge – for instance, by challenging the sufficiency of her evidence – does not undermine the sense that she knows; her evidence seems sufficient. However, in the HS context, her evidence does not seem sufficient, especially if it's challenged. For example, an interlocutor may raise the possibility that the bank has changed its hours. In LS this possibility is dismissed as not worth considering and, thus, fails to undermine her claim to know. But in HS, this possibility cannot be dismissed. Since the stakes are so high, even unlikely possibilities become relevant and bear on whether Mary can be said to know whether the bank is really open on the subsequent Saturday. DeRose argues that Mary does *not* know in HS. His explanation is that the truth conditions of attributions of the knowledge predicate shift with context. This is the heart of epistemic contextualism: whether the proposition expressed by "S knows that p" is true depends on the context.

For the most part I will not engage the large and complex literature that has developed around DeRose's examples and other similar ones. That literature turns on such details as the proposed causes of context-shift; for example, whether it is the epistemic context of the subject or the ascriber of the knowledge predicate. A number of positions have emerged responding to these cases. On the one hand, epistemic invariantists argue that whether the proposition expressed by "S knows that p" is true depends on the context.[21] There are many varieties of invariantism, including a moderate version according to which we can, irrespective of context, know many ordinary, quotidian propositions, but not propositions such as *I am*

---

[21]There's a good summary in Hawthorne (2004) and Turri (2010b).

not a brain-in-a-vat. On the other hand, many contextualists take the lesson to be that there is no single definition of knowledge, since whether it's correct to ascribe knowledge to an agent depends heavily on the epistemic context. However, I raise these debates only to set them aside.

The normative character of SRNA is invariantist because the necessary and sufficient conditions for warranted assertibility do not shift with context: it is always the case that an agent must satisfy both conditions (i) and (ii). And the locus of normative evaluation is always the speaker, not someone ascribing warranted assertibility: we judge whether an assertion is warranted based on whether the speaker, in the context, and based on her epistemic support, is warranted in making the assertion. We don't ask whether, based on the ascriber's epistemic position, the speaker is warranted.

The normative content of SRNA is contextualist because what it takes to satisfy (i) and (ii) will heavily depend on the context of assertion.[22] I take it as powerfully intuitive that whether one may warrantedly assert depends heavily on the context of assertion, both epistemically, conventionally, and pragmatically. Epistemically, we expect witnesses testifying in high stakes court cases to have especially strong evidence in order to warrantedly assert (discussed in Chapter 7), and we don't expect Mary asserting that the bank is open in the LS case to have especially strong evidence. The strength of reasons required to warrantedly assert thus seems to shift with the context of assertion. But the practice of assertion is even more complicated.

---

[22]I would further say that the context is fixed from the subject's perspective, not an ascriber's perspective. My position is thus not an ascriber-dependent contextualism. However, I am unsatisfied with the current selection of subject-sensitive positions to pledge allegiance to any currently on offer, such as Hawthorne (2004). It's also an open question whether the distinction that I'm making is wholly distinct from the epistemic contextualism debate, but I think that it is at least insofar as it's not clear that this distinction is being recognized in that debate.

Sometimes, as in Lackey's "selfless assertions," speakers may warrantedly assert even when they don't believe what they say. In some special contexts, then, other epistemic features such as belief aren't required, but in others they may be. In fact, there seem to be cases where one can warrantedly assert when one lacks epistemic justification: that is, one may assert what one believes or even knows to be false, apparently in violation of Grice's first sub-maxim of Quality.

Here is an example of the latter phenomenon. Suppose that Jim is teaching a Grade 6 science class. He wants to explain the structure of an atom and, more specifically, the electron configuration of different elements. Jim is well aware that an early model of the electron structure of atoms, the Bohr model, is no longer considered accurate. Under the Bohr model, electrons travel in restricted *orbits*. Electrons are treated specifically like particles, rather than more recent conceptions in quantum mechanics where electrons have both particle- and wave-like properties. The Bohr model often depicts electrons as planets orbiting the nucleus of an atom. There are also clear rules concerning how many electrons can be found in different orbits and it claims that electrons in different orbits do not interact. More recently, though, the Bohr model has been replaced with the *valence* model. Under the valence model, due to incorporating principles of quantum mechanics such as the Heisenberg Uncertainty Principle, electron "orbits" are replaced with probability "clouds." Electrons are not restricted to rigid orbitals; rather, they're patterns that represent the probability distribution of the position and angular momentum of the various electrons. Therefore, unlike the Bohr model, electrons in one orbital cloud can be found in different clouds. Strictly speaking, there's a non-zero

probability of finding an electron of an atom *anywhere.*

So knowing all of this, Jim also knows that his students aren't yet able to understand the valence model, but they are able to understand the Bohr model. So when it's time to teach his students about the electron structure of atoms, he asserts that "Electrons behave according to the Bohr model."[23] This is, strictly speaking, false and we may suppose that Jim knows that he's asserting something false.[24] However, he makes the assertion anyway because, pedagogically speaking, it is the best assertion he could make. Anything else, including a qualified assertion, would confuse his students or undermine the students' willingness to learn the concepts Jim wants them to.[25] I argue that Jim's assertion has supportive reasons: they're the institutional features of Jim being a teacher and having his students' learning needs in mind.[26]

One may object that Jim isn't actually asserting when he tells his students that electrons behave according to the Bohr model; he's possibly performing some other speech act. But this is implausible, I argue. In Chapter 2 I endorsed the definition of assertion as a speech act where one puts forward a proposition as both true and worthy of belief. An assertion licences the hearer to re-assert it and licences the hearer to use the proposition as a premise in inferences and practical reasoning. Furthermore, assertions create certain commitments on

---

[23]It's worth noting that this is how students are taught atomic physics: first students learn the Bohr model *as if it's true*, and are then taught that the Bohr model is incorrect. They're subsequently taught the valence model.

[24]One might object that while Jim's assertion isn't, strictly speaking, true, it's approximately true and this warrants his assertion. For example, if truth is the norm of assertion, it seems odd to criticize the assertion that it's 12:00 p.m., when it's really 12:01 p.m. While that's true for some assertions where precision isn't important, this objection doesn't apply to Jim's assertion. The Bohr model isn't even "approximately" true in the way that the assertion that it's 12:00 p.m. is sufficiently true for the assertoric context.

[25]Recall the discussion in Chapter 2 to the effect that obviously insincere assertions aren't taken seriously, though they're still *assertions*.

[26]Below I will argue that supportive reasons are still importantly epistemic.

behalf of speakers by requiring them, for example, to defend their assertions when challenged. All of these features are present in Jim's speech act in teaching his students.[27] He's putting forward the proposition that electrons behave according to the Bohr model as being true, worthy of his students' belief, and they can take that information and put it to use by re-asserting it, or by using it as a premise in forming inferences or practical reasoning. Moreover, Jim *does* feel a need to defend his assertion if challenged, a point I'll return to below.

So, I argue, Jim is making an assertion. The further question is whether it's warranted. I think that it is. Jim is properly respecting his students' epistemic needs *qua* learning. His asserting something false is done in service of later being able to teach his students the truth (as we understand it). But consider an enterprising student who learns through independent study that what Jim has said is, strictly speaking, false: electrons are thought to behave according to the valence model, not the Bohr model. Suppose that the student confronts Jim: "What you said is false. You lied to us! You shouldn't have said that." This is a *prima facie* apt challenge, one that requires a response from Jim. So, from Chapter 7, what would constitute a wholly adequate response? I think that Jim would likely appeal to his duties as a teacher and the epistemic licence it gives him, in cases such as this, to assert something that he knows is, strictly speaking, false.

Jim's supportive reasons include that such an assertion will provide his students with something worthy of belief (which may even rise to the level of justified belief), and that making this assertion is the best way for them to learn the attendant concepts. Jim's ped-agogical relationship provides supportive reasons for the assertion because, based on their

---

[27]And remember, adapting Rescorla (2009), these are constitutive *requirements* of assertion, not norms.

role as students being taught by a well trained teacher who has his students' learning needs in mind, they have good evidence for the truth of the proposition. That the students have these reasons is afforded by the structural and institutional justification imparted to Jim in virtue of being a good teacher and that asserting that p is the pedagogically best assertion he could make, and it's done with the aim of having his students later come to learn the truth. Moreover, the students can form a justified belief based on their (often implicit) understanding that good teachers tend to assert truthfully.[28] Were he to provide a qualified assertion, or skip directly to discussing valence model concepts, the students would struggle and Jim wouldn't be fulfilling his pedagogical duties. Therefore, the context of being a teacher, including the conventional and pragmatic features of the context, provides warrant for some assertions that a speaker knows to be false.

Insofar as Jim is making a *warranted* assertion, neither Lackey's RTBNA nor Douven's RCR can explain its propriety.[29]

RTBNA:

  i) One should assert that p only if it is reasonable for one to believe that p.

  ii) If one asserted that p, one would assert that p at least in part because it is

     reasonable for one to believe that p.

RCR: One should assert that p only if p is rationally credible to one.

---

[28]It's worth noting that Lackey (1999) uses this case as an argument against the transmission model of testimony: one can form a justified belief based on testimony without the speaker having a justified belief (or knowledge) herself.

[29]Lackey (2007); Douven (2006).

One might think that Jim is *lying* and think, with Lackey, that "telling a lie is a paradigmatic instance of an assertion that, though genuine, is in violation of the norms governing assertion."[30] So Lackey's view can accommodate Jim's speech act as an assertion. But since Jim is asserting something he knows to be false, it seems that he's lying and that the assertion is, therefore, unwarranted. While I leave open the possibility that there's a distinction between Jim's knowingly asserting something false and *lying*, let's assume that it's a lie. It's consistent with Lackey's view that, for overriding pragmatic reasons, Jim's assertion is warranted, but that it's not *epistemically* warranted. Since the norms of assertion under discussion are characteristically epistemic in nature, Lackey's view is committed to Jim's assertion not being epistemically warranted. However, if my argument to the effect that Jim's assertion *is* warranted all things considered, including in the epistemic sense, then Lackey's view leaves us without resources to explain this fact.

Since Jim knows that the assertion that p is false, the permissibility of the assertion would constitute a counter-example to Lackey's RTBNA. After all, it is not reasonable for *Jim* to believe the proposition and, therefore, the (non-existent) fact that it's reasonable for Jim to believe that p cannot justify his assertion. For similar reasons, the permissibility of Jim's assertion would constitute a counter-example to Douven's RCR. The problem, then, with both RTBNA and RCR is that these norms require that the *speaker* have the required epistemic reasons for believing that their assertion is true.[31] But while Jim judges that it's

---

[30]Lackey (2007, pp. 600-601).

[31]However, this characterization doesn't require that the speaker have a particular *doxastic* state in order to warrantedly assert. Lackey's view merely requires it be reasonable for one to believe the proposition asserted. A speaker may thus warrantedly assert without believing, or even justifiably believing, what she asserts.

not reasonable for him to believe what he asserts, it is reasonable for his students to believe the assertion. So the first condition of RTBNA could have been: one may assert that p only if it is reasonable for (the speaker or the hearer) to believe that p.

By "supportive reasons" I mean to focus on inherently epistemically inflected reasons for making an assertion. We know that speakers perform assertions for all sorts of reasons, just as someone may play poker for reasons other than strictly winning money. Consider the famous axe-murderer-at-the-door case: an axe-murderer comes and knocks on Franz's door, asking him if his friend, Ludwig, whom the murderer wants to kill, is in Franz's house.[32] Ludwig recently came, exasperated, asking for Franz's help hiding him, and Franz obliged. When the axe-murderer asks whether Ludwig is in Franz's house, Franz answers, "No, he's not here." He lied, but intuitions differ on whether this is a good thing to say in the context. Kant argued that one has a perfect duty not to lie: he'd say that this is an inappropriate assertion. A consequentialist, on the other hand, might argue that the assertion is morally justified since it maximizes utility to lie to the murderer (one supposes). The importance of this example is that *insofar* as one considers Franz's assertion morally appropriate, what seems to warrant the assertion is completely disconnected from epistemic considerations, and particularly the sort that I'm claiming constitute supportive reasons for assertions. So this may constitute a counter-example to my view.

Before responding to the challenge, consider a more recent example like one found in Brown (2008).[33] Steve is friends with Pete. Steve recently learned that Pete's wife, Shannon,

---

[32]Kant (1897).

[33]I'm altering the example considerably. In Brown's example, the husband (Pete) berates his friend (Steve) for knowing about the affair without telling Pete. The idea is that this case puts pressure on the view that

is having an affair. Pete, suspecting that his wife is cheating, asks Steve, "Is Shannon having an affair?" Not wanting to hurt his friend's feelings, Steve lies and says, "No, she's not." Now suppose, for the sake of argument, that Steve's assertion is warranted. Whatever reasons Steve has for making the assertion aren't epistemic: they're not with the aim of expressing knowledge or having Pete come to know the truth. Instead, Steve is actively trying to keep the truth from Pete. So like the axe-murderer case, insofar as Steve's assertion is warranted, it would constitute a counter-example to SRNA.

However, these cases are not counter-examples to SRNA. It's been a long-standing distinction in the norms of assertion literature that assertion is governed by a great many norms that are not *internal* to the practice. Since assertion is an action, it's governed by a variety of moral norms, some of which can *override* the norms internal to assertion, even constitutive norms like SRNA. In such cases, the assertion is appropriate, but not *qua* assertion. Franz is morally justified in his assertion, but not assertorically justified. So both Steve and Franz are assertorically unwarranted, but they're morally warranted in their respective lies. The moral norms override the assertoric norm.

I think that we can better understand why Jim, the science teacher, is *assertorically* warranted – due to the appropriate epistemic reasons – by contrasting Stella the creationist teacher with another version of a creationist teacher.[34] Stella asserts that, "*Homo sapiens* evolved from *Homo erectus*" even though she doesn't believe it. She recognizes that the scientific evidence supporting the proposition is overwhelming, but she doesn't believe it

knowing that p is sufficient epistemic support for asserting that p. The names are also my creation.

[34]The original example is from Lackey (1999) and Lackey (2007) and the second is an adaptation from Audi (2006).

due to a commitment to her faith and the claims of creationism. However, I've argued that Stella's assertion is warranted. Lackey's argument is that this demonstrates that one need not believe what one asserts in order to properly assert. She thus proposes RTBNA: Stella's assertion is reasonable for her to believe, and this warrants her assertion. I treat the case differently: it's *not* that it's reasonable for Stella to believe the proposition that warrants her assertion; rather, it's that her assertion has the aim of having her students come to believe what is best supported by the available evidence, despite her lack of belief. This is what we expect from good teachers. And notice the clearly epistemic nature of these reasons: she has the goal of putting forward a proposition as true in order for her students to come to believe something true (or even come to *know*).

The presence of these epistemic reasons is immediate in Stella's case since she's asserting something that, were the students to believe it, we plausibly think they would come to know it.[35] This isn't the case with Jim, since he's asserting something he knows to be false. However, his doing so is in service of the more distal goal of having his students come to later know how electrons behave in atoms by first learning something false (though they don't know that this is what's happening until later). Jim's performing his assertion in service of the further epistemic goal is what provides him warrant to assert something he knows to be false. If he were doing it for some other reason, such as his joy in asserting false things to his students in order to trick them, then these reasons would not provide *supportive* reasons (and hence warrant) for his assertion. To see this more clearly, consider the following parallel to Lackey's creationist teacher.

---

[35]This is the thrust of Lackey's use of the case against the transmission model of testimony.

230

Admittedly, Lackey's description of Stella isn't entirely plausible: what dyed-in-the-wool creationist believes that the evidence in favour of evolution is overwhelming? So suppose that there's a more plausible creationist teacher, Norbert, who's tasked to teach evolution, who not only doesn't believe in evolution, but he thinks that there's no good evidence in its favour: the scientists are all deluded.[36] But he has to teach the same curriculum as Stella, and so he also asserts to his students, "*Homo sapiens* evolved from *Homo erectus.*" So the context of Norbert's assertion is identical to Stella's, and the proposition asserted is also identical. It follows from SRNA, then, that the quality of supportive reasons required to warrant Norbert's assertion will be the same as for Stella's. The difference is that the *only* reason that Norbert asserts what he does is that he'll be fired if he asserts otherwise. That is, he just wants to keep his job: he doesn't care about his students' learning needs in this case. If the School Board were to suddenly allow him to teach creationism instead, he would. So in a sense he's a *mercenary* teacher of evolution.

Stella's assertion is warranted, but Norbert's is not. Both speakers occupy identical contexts of assertion: they're teachers teaching something that they don't believe. Moreover, the content of the assertions (the proposition) is the same. The only difference is that the reasons for Stella's assertion are the right sort: they're properly connected to the epistemic nature of good assertions. Good (*i.e.* warranted) assertions are those that properly aim at expressing knowledge. Lackey's correct that one need not always believe what one asserts, but her proposed norm cannot explain the impropriety of Norbert's assertion, whereas SRNA can. Norbert doesn't have good (*i.e.* supportive) reasons for his assertion: he's not aiming

---

[36]As noted above, the example is an adaptation from Audi (2006).

at expressing knowledge.

These cases should help illuminate, then, why Jim is warranted in his assertion (of something he knows to be false) and how what warrants his assertion is inherently epistemic. Although in the immediate context of the assertion he's not aiming at expressing knowledge, in the wider context of teaching concepts in a progression (from easier to more difficult, and from the false to the true) he is aiming at expressing knowledge. Warranted assertions are those that properly aim at expressing knowledge *in the context* of assertion. In Jim's context, he has the appropriate supportive (*i.e.*, epistemic) reasons for his assertion. If Jim were to make the same assertion, in the same context, without the aim of students coming to know the valence model (or without situating the assertion as part of a progression), then he wouldn't be warranted.

Since SRNA is highly sensitive to context, one might wonder if it's a constitutive norm. We saw in Chapter 3 that many authors, following Timothy Williamson's lead, have attempted to identify a "simple" constitutive norm of assertion. A norm is simple if it takes the form of a C rule (One must: assert p only if p has C) and "all other norms for assertion are the joint outcome of the C rule and considerations not specific to assertion."[37] But why think that assertion is governed by a simple norm, let alone a norm both simple enough to accommodate a C rule, but substantive enough to support the proviso that all *other* considerations relevant to assertion will not be "specific to assertion"?[38] Assertion is a fundamentally social practice, after all. So at first blush it shouldn't be surprising if

---

[37]Williamson (2000, p. 241).
[38]Both Shieber (2009) and Rescorla (2009), and probably others, raise this point.

there were social features specific to or internal to assertion. In fact, I argue below that there *are* conventional and pragmatic elements specific to assertion and that these elements should be built into the norms of assertion. One consequence of this is that assertion is not governed by a simple C rule; another is that satisfying the necessary and jointly sufficient conditions for warranted assertibility varies with and conventional and pragmatic contexts.[39] However, SRNA is a constitutive norm: in order for an assertion to count as assertorically warranted, the act must satisfy SRNA, although some assertions may be morally appropriate (but assertorically inappropriate) such as some noble lies.

A distinctive feature of SRNA is that it directly incorporates conventional and pragmatic features of context into the norm, rather than considering these elements as in force but external to the practice. It is widely recognized that there are important conventional aspects to assertions and that these may bear on assertoric norms. Although I might know that my boss's breath smells bad, it would be impolite to assert, "You have bad breath" in the middle of a department meeting. The conventional aspects of being in a meeting with my boss make the assertion inappropriate, even if I know the proposition asserted. Few philosophers think that it follows from KNA that the knowledge rule (assert only if one knows) is a sufficient condition for warranted assertibility.[40]

Moreover, there are specialized linguistic communities that may create their own conventions for proper assertion. It's plausible, for example, that a group of mathematicians might

---

[39]In fact, I'm open to further context shifts due to speech act contextualism. See Turri (2010b).

[40]Though there is now a discussion whether knowledge is the sufficient *epistemic* condition for assert. See Lackey (2010) and Carter & Gordon (2011). SRNA would say that *of course* knowledge isn't sufficient *simpliciter* because such a claim ignores the context of assertion. In some contexts knowledge will be the sufficient epistemic condition, but in others it won't, and we should expect this.

gather at a conference and agree upon a rule of assertion: one may assert only if one can prove the truth of one's assertion. This isn't meant as a norm that will govern assertion generally; it's simply a fiat norm created for their own immediate and local purposes. As an example that may be more familiar, professors cannot accuse a student of cheating, even if they know that the student cheated, unless they can provide proof.

It is also often recognized that there are important pragmatic aspects to assertions that may bear on assertoric norms. There are at least two senses in which we find pragmatics in assertions. The first involves the role of assertions as reasons for action. Suppose that we are at a train station (discussed in Chapter 3) and I think that train #3 is your train. I have some evidence for this, though I recognize that I don't know because I have my doubts. It may be permissible for me to assert, "That's your train, run!" if the consequences of missing the train are dire, but the consequences of my being incorrect are not dire.

The second sense involves cases where a speaker has options in what to assert and what is asserted is, strictly speaking, known but misleading to the hearer. Suppose that I just saw Jim walk into his office. Mary comes to me and asks "Is Jim in his office?" I respond, "It's possible." Since I know that Jim is in his office, it's true (and also known) that it's possible for Jim to be in his office. However, this assertion carries a misleading implicature of doubt. Although I know that it's possible for Jim to be in his office, the assertion is not warranted due to the pragmatic implicature.

In such cases conventional and pragmatic features of an assertion may appear to "override" the normal conditions for warranted assertibility. However, I think that it is inapt

to describe such conventional and pragmatic features as "overriding" some norm; rather, these features contribute to fixing the assertoric context. Since context partially determines what is required for warrantedly asserting, I suggest a norm that incorporates these features into the norm, rather than attaching them as an external addendum. However, there may be some conventional and pragmatic features relevant to the appropriateness of assertion which are external to the practice and, in this sense, would override the norm (*e.g.*, moral considerations).

The Williamsonian position is to consider these pragmatic features, not as part of the practice of assertion, but as more general conversational norms.[41] But some of these features are not in place for other speech acts such as conjecturing. Or, at least, if there are prudential features to conjecturing, they often differ from the prudential considerations present in asserting. It's hard to be rude or misleading, for example, when conjecturing.[42] (Though one can certainly be rude or misleading when questioning.) Although "be polite" is a general conversational rule, there are conventional aspects that are unique to assertion. It is a conventional feature of assertions, *e.g.*, that a speaker be willing to defend an assertion in the face of a challenge, or else retract the assertion.[43] This isn't present in questioning, conjecturing, or commanding. In the face of a student challenging Jim's assertion based on evidence against the Bohr model of electrons, in favour of the valence model, Jim is required by conventional features of assertion to respond in a particular way; namely, defending his

---

[41]Rescorla (2009) also takes this position. It's a fairly common response.

[42]However, there are some rude conjectures. For example, "I conjecture that you're a jerk." We can notice, though, that rude conjectures seem slightly more artificial than rude assertions such as, "You're a jerk."

[43]Compare Kvanvig (2009) and Rescorla (2009), for example.

assertion, retracting it, or perhaps modifying it (*e.g.*, replacing it). So why consider these as not essential to constitutive norms of assertion *qua* doing assertion properly?

The motivation seems to arise from a scope error. We recognize that there are conventional and pragmatic features that, at one level of description, are common to a number of speech acts. Grice's Cooperative Principle, for example, seems to apply to all speech acts.[44] But it doesn't follow that there aren't conventional or pragmatic features of politeness specific to assertion. In fact, I think that there are such features specific to assertion, particularly when we consider the special role assertion has in our communication. There are particular conventional burdens on asserters, such as that one shouldn't misrepresent oneself. Moreover, the conventional commitments taken on by a speaker in asserting are much weaker (or non-existent) in other speech acts, such as conjecturing. Moreover, since we characteristically use assertions as premises in forming inferences or in our practical reasoning, assertion has pragmatic features that other speech acts will not. That is, it's a defining characteristic of assertion that it has these conventional and pragmatic features. Since so many conventional and pragmatic features seem internal to the practice of assertion, in addition to other more general external features common to many speech acts, I see no reason to keep such assertion-specific conventional and pragmatic features of proper assertion separate. I therefore include them into my proposed central norm of assertion, SRNA.

SRNA, then, is a norm *of* assertion. That is, it's a norm internal to the practice and not merely an external, more general norm. The pragmatic and conventional features in-

---

[44]Except, for example, jokes or insulting, provided that insulting is a separate speech act. It's a feature of some jokes that they explicitly violate Grice's norms.

corporated to the norm are internal aspects specific to assertion. Insofar as more general conventional and pragmatic features (and norms) govern assertion, then I follow Williamson (and Rescorla) that such features are not norms *of* assertion and so do not need to be incorporated into the norm. The mistake, I argue, lies in not recognizing or giving too little weight to the fact that there are at least some such features internal to the practice.

8.5 Performance Normativity and Meta-Aptness

Sections 8.2 to 8.4 present the basic outline for the case for SRNA. In this section I discuss a recent defence of KNA in normative terms that are somewhat similar to how I conceive of SRNA. Specifically, Ernest Sosa considers assertion to be governed by a norm conceived in terms of *performance* normativity: assertion should be assessed much like we assess performances such as an archer shooting at a target. If this is how we ought to conceive of norms of assertion, then we expect to see norms of assertion analogous to normative assessments of actions like the archer's shot. Sosa conceives of performance norms by considering their relationship to the (constitutive) goals of a practice. How we evaluate a performance depends on what goals an agent has. We don't necessarily assess the archer's shot merely by considering whether it hit its target because, for example, maybe the archer wasn't attempting to hit the target: perhaps she had a different goal (*e.g.*, it was meant as a warning shot: intended to be close to, but ultimately missing, its target). Both of these concepts are central to how I conceive of SRNA, but Sosa uses them to arrive at KNA rather than a reasons-based norm. So his view is a useful foil for my position, and discussing our differences will draw out some

important further features of my view.

One of Sosa's most important contributions to epistemology is conceiving of knowledge as an action. He argues that knowledge has a "triple-A" structure: knowledge is *accurate*, *adroit*, and *apt* belief. A belief is accurate when it is true, adroit when a belief is competently cultivated, and apt when a belief's being true manifests an agent's competence in believing it. Knowledge on one level ("animal knowledge") is apt belief. Higher-order or *reflective* knowledge comprises an apt belief that is also meta-apt: that is, knowledge generated when an agent competently evaluates whether to form a belief (*i.e.*, competently evaluates risk). Finally, a belief is fully apt (and an agent knows full well) when "its first order aptness derives sufficiently from the agent's assessment, albeit implicit, of his chances of success (and, correlatively, of the risk of failure)."[45] That is, a belief is fully apt when a belief's aptness sufficiently manifests its meta-aptness.

Sosa's favourite illustrative example is of an archer shooting at a target. If we merely evaluate an archer's shot based on whether a particular shot hit a particular target, then we're forced to regard the successful shot of a lucky novice as normatively indistinguishable from an equally successful shot by a professional archer. Analogously, if all we cared about belief was that it successfully attains truth, then lucky guesses would be normatively equal to an expert's coming to believe the truth of a proposition through the exercise of expert knowledge and methods.[46] To forestall this implausible consequence, we add a competence requirement: true beliefs must be competently obtained.

---

[45]Sosa (2011, p. 11).

[46]One might object that we often don't believe our guesses. However, there are enough cases where people do form beliefs without adequate evidence, and we call these guesses.

Next consider an expert archer's shot that is properly aimed at a target and, unsurprisingly, hits its target.[47] But imagine that during the arrow's flight, a gust of wind redirects the arrow off course, followed by a second gust of wind which compensates for the first gust, putting the arrow back on target (and subsequently hitting the target). In such cases, we'd say that the shot is a competent one: were it not for any wind, the arrow would hit its target. But, on Sosa's account, the shot's accuracy doesn't *manifest* the archer's competence. The archer's shot is analogous to Gettier beliefs: justified true beliefs, the truth of which do not manifest the agent's abilities. So Sosa argues that knowledge, as with any performance normativity, requires for its full success that the success manifests the agent's competence: an action, as with belief, must be apt.

Sosa argues that knowledge is normatively related to action in a means-ends intentional relationship between belief and truth. That is, in believing that p, through some belief forming process (the means), we intend to believe the truth (the end).[48] Suppose, by way of analogy, that an agent intends to turn on a light. She knows that the best way to do this is to flip the light switch. All performances that aim at a goal, Sosa argues, fall under the AAA structure. In order to be competent in performing an action (and especially in order to be meta-apt), an agent must form the relevant means-ends belief.[49] Here, for example, the agent must form the belief that flipping the switch will bring about her aim of turning

---

[47]What's the difference between aiming and *properly* aiming? One can aim to hit an archery bullseye by closing one's eyes, spinning around five times, and shooting. But this is a poor method of aiming: it's not exercising one's skills in aiming at the target. One properly aims at a target when one exercises a skill in aiming to achieve a goal.

[48]Sosa (2011, pp. 44-46).

[49]I'm simply explicating, not endorsing, Sosa's view. If readers find this aspect of his view implausible, then all the better for my arguments.

on the light. Suppose that she flips the light switch and the light turns on, but her action is Gettiered: there's a short in the circuit, but, luckily, a mouse has left a piece of tin foil next to the wiring, so that the circuit closes properly and the light turns on. Sosa argues that the agent's means-end belief is not apt since its accuracy doesn't manifest her competence (in selecting the action of flipping the switch). Furthermore, the performance (the flipping the switch) fails to be apt. "It attains success by luck, in a way that is relevantly deplorable."[50] So the agent's action is worthy of (normative) criticism.[51]

The critical point is that "any action taken as a means to a further objective will of course fall short if it does not bring about that further objective."[52] Sosa argues that,

> [i]napt performances fall short not only in that they might have been better on relevant dimensions. They fall short in the fuller sense that they fail to meet minimum standards for performances. Because they are inapt, they are therefore flawed: not just improvable but defective.

So any Gettiered performance, or any inapt performance, is normatively defective; not merely in that such performances could have been better, but to the point of being deplorable. However, as I discuss below, this view mistakes performing an action that turns out wrong (not as one wanted – *i.e.*, that falls short of the goal) for performing an action *wrongly*.

As mentioned previously, Sosa also conceives of assertion as an action and, consequently, something that falls under the AAA structure of performance normativity. He supposes,

---

[50]Sosa (2011, p. 45).

[51]As opposed to, say, moral criticism. There are different brands of criticism. One can fail to obtain knowledge, and hence be worthy of epistemic criticism, while still being morally praiseworthy, and vice versa.

[52]Sosa (2011, p.46).

quite plausibly, that sincerity is a norm of assertion: one ought to assert only what one believes, and that one ought not to assert what one doesn't believe. In Chapter 2 we saw that this is sometimes viewed as a necessary condition for assertion (*i.e.*, Searle) or at least a necessary condition for warranted assertion (*i.e.*, Bach and Harnish). Moreover, we saw in Chapter 6 that the Moorean view of assertion is that when one asserts that p, one represents oneself as (at least) believing that p. But, in light of this, what are we to make of cases such as Jim the science teacher, or Lackey's creationist teacher example? Sosa suggests that such cases can be accommodated by making a distinction between making an assertion "in one's own person" and "as occupier of a role." A teacher only asserts (insincerely) properly when viewed as asserting *qua* teacher but not *qua* (regular) person. After all,

> As a newscaster or as a teacher one may be called upon to say things, and thereby to assert them, as in the classroom or in a newscast, even when one does not believe what one says. To play one's epistemic role in such contexts may just require reading (assertively) from a script, or from a teleprompter, or reporting from memory, where one serves as a mouthpiece for a deeper intuitional source of the information conveyed, the deeper source that is the school, or the news organization. So I will assume that sincerity is a norm of assertion in one's own person, where one is not playing a role in some epistemic institution (for the delivery of information or the like).[53]

One problem with this view is that it seems to require all sorts of specialized norms of

---

[53]Sosa (2011, p. 47).

assertion, depending on one's institutional role. Shall we propose a separate norm of assertion for doctors, teachers, plumbers, and ice cream vendors? It would surely be more explanatorily satisfying to have a norm that can unify our intuitions about how the conditions for assertibility shift with the context of "asserting in one's own person" compared to "asserting as occupier of a role." This is what SRNA purports to do: it explains why a teacher, in the context of teaching, given the conventions and pragmatics of the situation, is permitted to make some insincere assertions. Of course, this carries the caveat that the relevant conventional and pragmatic features are present. Not just any insincere assertion is licensed by being a teacher or a newscaster. And SRNA gives us an explanation for why some insincere assertions are warranted but others are not.

There is another serious problem with Sosa's account. He begins his analysis of performance normativity by considering plausible necessary and jointly sufficient conditions for knowledge. Since it's widely accepted that knowledge is factive (*i.e.*, truth is a necessary condition), it makes sense to include an accuracy condition. This makes truth, or obtaining the goal of an act of believing, a necessary condition for propriety: anything less is defective, he says. But while this is plausible for an analysis of knowledge, it simply begs the question against non-factive norms such as the performance normativity of other actions such as assertion. The very question we're concerned with is whether the central norm of assertion is factive (*i.e.*, KNA, CNA, or TNA) or whether a reasons-based norm is adequate. One shouldn't merely assume that what applies to performance normativity for the construction or attainment of knowledge applies to performance normativity for assertion, or for actions

more generally. To assume that accuracy is a necessary condition for all apt action is to beg the question. While it's plausible for knowledge, further argument would be required in order to apply it to other actions such as assertion or practical reasoning.[54]

To be fair, though, Sosa does think that some support for his view that accuracy is required for apt action (*i.e.*, appropriate action) comes from his discussion of the means-ends relationship between an action and its goal. Doesn't an archer do something wrong, after all, when she misses her target, even when it's through no fault of her own? Perhaps she's blameless, but still worthy of criticism (in a sense cashed out in terms of primary/secondary propriety, for example). But we have already seen that this is a confused and unilluminating way of describing such situations. This is precisely the sort of case in which we must clearly distinguish our evaluations of actions from our evaluations of outcomes; phrases like "did something wrong" are just too blunt an instrument for these purposes.

A brief reminder of our earlier discussion of norms may be useful here. Suppose that I offer you the following bet: I'll bet you your $100 to my $1,000,000 on the roll of a one hundred-sided die. If it turns up #1, I win, anything else, you win. So you have a 99% probability for winning this bet, which has a very high expected value. Assuming that you can afford to lose the $100, you would be rational to take the bet (and possibly irrational to pass). But suppose that you lose anyhow, as will happen in 1% of such cases over the long run: shall we say that you acted wrongly in taking the bet merely on the grounds that you lost? I suggest that we shouldn't think this way – indeed, that it's obvious that we shouldn't think this way. You

---

[54]And while one may object that this is merely a dialectical objection, the objection cuts at the heart of Sosa's argument for the knowledge norm of assertion.

made a great decision; the outcome is irrelevant to our normative evaluation of your action, except inasmuch as it might lead us to revise our confidence in the statement of probabilities. Of course, you were aiming at winning – that was the goal, the whole point of taking the bet – but it fails to follow that, simply because you failed to achieve your goal, your action failed to satisfy the normative requirements for the action, let alone the "minimum requirements" as Sosa puts it.

By the same token, if an expert archer properly aims and fires at her target, but a bird flies in front of her arrow just as it is about to hit the bulls-eye, we don't say that the archer did anything wrong in her action.[55] We do say that she failed insofar as she failed to obtain her goal, but this is an evaluation of the outcome of her action, not a normative assessment of the performance itself. That is, it's a *description* of the agent's action rather than a normative assessment of it. Since we're concerned with performance normativity, not outcome normativity, we evaluate the performance (the action). And, as I have argued throughout, the results aren't relevant to the normative evaluation of an action. To say of the unsuccessful die bet that it fell short *qua* performance normativity is to commit a common fallacy found in gambling. (The fallacy is colloquially referred to as being overly results-oriented.)

One might object, though, that the results of an action cannot be, strictly speaking, irrelevant to the normative assessment of an action. In a sense, this is true. In Chapter 3 I argued that norms derive their content, in part, from what it means to properly (*i.e.*,

---

[55]Compare this with the infamous Randy Johnson pitch that killed a bird in an MLB baseball game on March 24th, 2001.

rationally) aim at achieving the goal of a practice. We take the die bet because, in the vast majority of times we take the bet, we win. Modally speaking, even when we lose in the actual world, we win in most nearby possible worlds. That we obtain our goal in most nearby possible worlds is one way of cashing out an action being a good means for achieving the goal of a practice. But it's only in this very weak sense that outcomes are relevant to the normative assessment of an action: provided that the action is one that properly aims at achieving the goal of a practice, then whether it actually achieves the goal (*i.e.*, the outcome) is irrelevant to the normative assessment of the performance.

Suppose that instead of taking the die bet yourself, you're the manager of an investment account. You're charged with investing a portion of a client's life savings. You're presented with the forced choice to invest 10% of the account's balance (let's say it's $10,000) on the roll of a one hundred-sided die. If you win, I'll pay out 100-to-1 ($1 million). You win on any roll other than #1. This bet carries a very high positive expected value. If we suppose that it's rational to take the bet (ignoring loss and risk aversion) how do we think that your client will view your taking the bet and losing? I think that the client would be pleased with you: sure you lost, but you made a great investment. They may say, "You win some, you lose some: it's not your fault that we lost." We would criticize someone for instead saying, "Why did you take that bet? You lost!" We'd accuse them of not knowing how investing works: we take calculated risks and "let the chips fall where they may." The bet had a 99% chance of succeeding with a very high expected value. But suppose that instead of accepting the bet on these terms, you prefer to bet on the roll coming #1: so you pay $10,000 on any

245

roll other than #1, and I pay $1,000,000 on only #1. This is essentially an even-money bet: you're getting the right odds for the probability of winning such that the expected value is neutral: you don't stand to win money in the long run – the probability and prize of winning is balanced by the probability and cost of losing. Let's suppose that you win: how will the client view your choice of the alternate wager, even though you win? I think that the client would reproach you for such an absurdly risky bet (with no expected value) while turning down a highly profitable (relatively) safe bet. A risky but successful bet is viewed as a worse action than a safe but unsuccessful bet. Sure, *ceteris paribus*, a successful bet is a better outcome than an unsuccessful one, but these are out of our control. We normatively evaluate actions, generally, based on what we can control: our actions, *not* outcomes.

We can draw evidence for this view from all sorts of observations of everyday interactions. Sports commentators make this error almost daily. Suppose that you're watching a close baseball game. It's in the bottom of the 9th inning and your team is up to bat down by one run with two outs. There's a runner on first and the league-leading home run slugger is up to bat. He's having a great day, hitting 4/4 with a home run. Now suppose that the manager decides to pinch hit for the slugger, putting in an unproven rookie, "on a hunch." It should be obvious that this is an extremely risky play, with a low probability of success. But suppose that the gamble pays off and the rookie, who's never hit a home run in his life, and doesn't even have a big league hit, hits a home run and wins the game. We can imagine a commentator saying, "I thought the manager was crazy when he made that substitution, but it looks like he made the right decision after all!" We should cringe: this is terrible reasoning

on the part of the commentator. We shouldn't retroactively assess a decision based on the results: taking a high risk, low probability, low expected value bet is a bad decision when a low risk, high probability, high expected value bet is also available (*i.e.*, the slugger!). The moral is that we assess the decision at the time of decision (partly because results are irrelevant).

In making this point more fully, suppose that the manager went with the slugger and he struck out thus losing the game. And suppose that, had the manager gone with the rookie, they would have won. My argument is that even in cases such as this, it's incorrect to say that the decision to go with the slugger was wrong. Sure, the outcome wasn't what the team wanted, but the evaluation of the manager's decision doesn't depend on the outcome: it depends on whether the decision, for example, maximized the team's chances of winning. Going with the high probability, high expected value decision does this; going with the rookie does not.

It's important to recognize, though, that one particularly appealing aspect of Sosa's view is the connection between aptness, meta-aptness, and the gradability of performance normativity. For Sosa, an agent knows if she has apt belief. But an agent has higher-order knowledge, reflective knowledge, when the aptness of her belief manifests a meta-aptness: a competence in evaluating the risk of initiating a performance. The expert huntress has to decide not just how to execute a shot (aptness), but whether to take a shot – say, because arrows are costly, or because missing may scare the prey. The competent selection of a shot based on taking these features into account is meta-aptness. And if the shot is successful in

part because the shot was so selected, then the aptness manifests the agent's meta-aptness, and the action is fully apt. So someone taking a bet, who also takes into consideration the relevant risk factors (whether there's too much risk-of-ruin, whether there are better bets available), is better than a similar agent who takes the bet without taking these factors into consideration: a fully apt action is better than an apt action. However, performance normativity lies at the level of aptness: fully apt actions are not required for appropriate action, they're just better. Aptness is what determines the minimum requirements for appropriate action.[56]

We can apply this to assertion. Suppose that SRNA is the central norm of assertion. And suppose that for the context, a speaker merely needs to have a justified belief: asserting that p will be warranted only if the agent is justified in believing that p. Since the agent satisfies these conditions, the assertion that p is warranted. But suppose that the agent can do one better: the agent actually knows that p. We can say that this assertion is better than one based only on justified belief without therefore thinking that the better assertion is the standard for warranted assertibility. This is an error common to KNA and CNA advocates, I submit. They recognize that *ceteris paribus* a true, known, or even certain assertion is better than one merely based on justified belief, but they infer that the "better" epistemic state is the norm.[57] This mistake is a version of mistaking the goal of assertion with its norm.

So what might be the goal of assertion? It's quite plausible that assertion aims at ex-

---

[56]It's also important to note, though, that my use of "aptness" here is slightly different from Sosa's in that I don't think that actions such as assertion necessarily have factive norms. So my use of aptness doesn't include an accuracy condition in the way that his does.

[57]And, moreover, the lower quality action is "criticizable" for not being of higher quality.

pressing knowledge. In fact, I think that the constitutive aim of assertion is to assert known propositions. CNA advocates sometimes argue that, *ceteris paribus*, assertions expressing certainty are better than assertions expressing knowledge. Therefore, they argue, certainty is the norm of assertion, not knowledge. KNA advocates concede the point that, *ceteris paribus*, assertions expressing certainty are better than assertions expressing knowledge (how could one not?). But KNA advocates argue that it doesn't follow that certainty is the norm of assertion. An assertion expressing a stronger epistemic state being better, *ceteris paribus*, than an assertion expressing a weaker epistemic state does not establish that the former is the norm of assertion. Norms allow for gradable evaluations, after all. It's possible for one action to satisfy a norm "better" than another action where both actions satisfy the norm. By analogy, if an archer's goal is to hit the bulls-eye of a target, we might praise one archer over another for more masterfully achieving the goal even if both archers hit the target according to the norm of archery.

It's important that any proposed norm of assertion respect the intuitive gradability of norms. While an assertion expressing a weaker epistemic state than knowledge may satisfy the norm, a similar assertion that expresses knowledge will be, *ceteris paribus*, better. SNRA respects this distinction. My proposal, then, is that while meta-apt performances are better than merely apt performances, meta-aptness isn't required for proper action.

8.6 Conclusion

In this chapter I have presented my proposal for the central norm of assertion: SRNA. It's a reasons-based norm that directly incorporates the pragmatic and conventional features of assertion, rather than viewing them as external to the practice. Since we're concerned with norms of assertion, we're concerned with normative features internal to the practice. This extends to pragmatic and conventional features specific to assertion; the norm should therefore reflect this. I have modelled SRNA on how we understand the relationship between the norms and goals of betting. Based on this model I have argued that, in terms of performance normativity, the results of an action are importantly irrelevant to the normative assessment of an action. Outcomes are only relevant in that norms dictate what it means for an action to properly aim at achieving the goal. SRNA also respects our intuitions that performance normativity allows for gradable evaluations: some actions are better than others where both actions are appropriate. Finally, I have used Sosa's recent discussions of performance normativity, and his means-ends view, as a useful foil for my own.

In Chapter 9 I will take stock of the various arguments for and against various norms of assertion with the aim of explaining how SRNA measures up against its competitors in accommodating the range of data canvassed in this dissertation. I argue that SRNA can explain each of the individual classes of data, such as the lottery paradox, Moore's paradox, and the prompts and challenges data, at least as well as any other proposed norm, and often better, and can best explain the data when taken as a whole. Furthermore, SRNA has the advantage that, insofar as it does not aspire to complete generality, its commitments are less

counterintuitive or problematic than other views – notably KNA.

Chapter 9

Conclusion

In Chapter 2 I discussed how we should understand the characteristic features of the speech act of assertion. I argued that assertion involves a speaker putting forward a proposition as true, and perhaps as worthy of belief. Since assertion is characteristically a social activity, speakers take on certain commitments in terms of a duty to defend one's assertion. Assertions give hearers a *prima facie* reason to believe what's asserted, which connects with contemporary discussions in the epistemology of testimony.[1] Since assertion is an important method for communicating, there is a strong intuition that there are right and wrong ways to go about asserting. This is why we think that there are norms governing assertion. The principal aim of this dissertation is to make a case for a reasons-based norm of assertion. In this chapter I take stock regarding how well my proposed norm SRNA can explain the range of data often used to support other norms such as TNA and KNA.

---

[1]See, for example, Lackey & Sosa (2006).

9.1 The Linguistic Data and Norms of Assertion

In Chapters 5 through Chapter 7 I discussed a variety of linguistic data most often offered in support of KNA: the lottery paradox, Moore's Paradox, and challenges to assertions, respectively. It is sometimes argued that KNA can, on balance, best explain each of these sources of data. What best explains the intuitive unassertibility of, "Your ticket didn't win," is that one cannot know that the ticket didn't win and that knowledge is the norm of assertion. What best explains the intuitive unassertibility of "Dogs bark, but I don't know that dogs bark," is that in asserting that p one represents oneself as knowing p, so in asserting, "Dogs bark, but I don't know that dogs bark," one explicitly denies in the second conjunct what one represents oneself as knowing in the first conjunct, and knowledge being the norm of assertion explains this. And what explains the near universal propriety of challenging an assertion with, "How do you know?" is that such challenges challenge a speaker's knowledge, which implies that one should assert that p only if one knows that p, which is KNA. So not only does KNA well explain each of these sources of data when considered individually, but it provides a unifying analysis of the data when taken as a whole. So KNA seems very well supported by the data, at least on the face of it.

One of my aims with this dissertation has been to knock down each of these three pillars of support for KNA, in each instance offering an alternative explanation by a reasons-based norm. In Chapter 5 I offered an analysis for why one cannot know lottery propositions such as, "Your ticket didn't win," when made prior to information on the results of the draw. I employed a relevant alternatives theory (RAT) which I called the *Non-Destabalizing*

Alternatives Theory (NDAT). The important feature of lottery propositions is that they remain unknowable and unassertible even though they may be true. I argued that since one can truly believe a lottery proposition, absent of Gettier conditions, the only explanation for the unknowability and unassertibility is that one lacks sufficient (epistemic) *justification* for the proposition. The upshot is that one need not invoke a knowledge norm of assertion to explain the unassertibility: any norm which includes a justification requirement (or what I've simply called *reasons*) will be able to explain the phenomena. The fact of the matter is that one *never* has the supportive reasons for lottery propositions in order to warrantedly assert them. Thus SNRA properly predicts the lottery paradox data.

Part of my argument here for SRNA is that there are good reasons to think that a reasons-based norm is *better* suited to explain the linguistic data than KNA. If what seems to explain the unknowability and unassertibility of lottery propositions is that agents lack sufficient (epistemic) justification, then invoking a knowledge norm is overkill: all we need is a reasons-based norm. Both a nut-cracker and a sledge-hammer will successfully crack a walnut, but we only need the former. Moreover, the former seems best-suited for the task. I think the same is true for SRNA *viz.* the lottery paradox data. KNA's explanation depends on the lack of adequate epistemic reasons for warrantedly asserting (or knowing), so a lack of reasons is at the heart of the impropriety of asserting a lottery proposition. Therefore, SRNA explains the data at least as well as KNA *and* enjoys a theoretical advantage in rendering the KNA account superfluous

In Chapter 6 I discussed Moore's Paradox and the intuitive unassertibility of sentences

like, "Dogs bark, but I don't know that they do." G.E. Moore's solution is to treat a speaker's asserting that p as the speaker's representing herself as believing that p. In fact, he argues that in asserting that p, one represents oneself as *knowing* that p. I disagree with this characterization of assertion. It's not clear that in asserting a proposition, a speaker represents herself in a particular mental state (whether belief or knowledge). As I've argued throughout, there are all sorts of warranted assertions where speakers have varying mental states. Jim the science teacher (from Chapter 8) warrantedly asserts something he knows to be false, whereas Stella (from Lackey's creationist teacher example) warrantedly asserts something she doesn't believe, but something for which she has epistemic justification. I've adopted a view of assertion which minimally characterizes it as a speech act where a speaker puts forth a proposition as true, and as offering reasons for the hearer to believe what's asserted. But putting forward a proposition as true, and perhaps worthy of belief, is very different from asserting necessarily representing the hearer as having a particular mental state. At best, I think, in asserting that p one represents oneself as being warranted in making the assertion. Alternatively, one could say that in asserting that p, one represents oneself as having *assertoric authority* to assert that p. Perhaps it's plausible that when testifying[2] or telling[3], one represents oneself as being in a particular mental or epistemic state (knowing, being justified, justifiably believing, or truly believing), but I don't think that this is the case for assertion.

Returning to the question of explaining the apparent impropriety of asserting Moorean sentences, however, Igor Douven argues that a reasons-based norm can equally well explain

---

[2]See, for example, Goldberg (2007).
[3]See, for example, Moran (2006).

the Moorean data as KNA can.[4] Let's assume that SRNA is the norm of assertion. In that case, one should assert that p only if one has sufficient epistemic reasons for the context. If we assume a standard context, would one generally have good reasons to put forward a proposition like, "Dogs bark, but I don't know that they do"? Probably not. One would have to have good reasons to think that both conjuncts are true: that *Dogs bark* and that *I don't know that dogs bark* are both true. It will rarely, if ever, be the case that a speaker has adequate epistemic reasons to think that both conjuncts are true: usually, when we have adequate evidence to assert the second conjunct, these reasons *defeat* our reasons for asserting the first conjunct. So SRNA properly predicts this data.

In fact, there are potentially all sorts of propositions that are Moorean but warrantedly assertible in the right contexts. Douven (2006) argues that one may have good evidence for *That's a barn* but also have some evidence that one is in a Gettier situation. Since one may be epistemically justified in believing that p even in Gettier cases, one could have good reasons to think that both *That's a barn* and *I don't know that that's a barn* are true. So provided that the second condition of SRNA is satisfied, there are cases where one may warrantedly assert Moorean propositions. Other potential cases of warrantedly assertible Moorean propositions arise from special contexts like Jennifer Lackey's examples of selfless assertions.[5] Since one can fail to know that p because one fails to believe that p, one may assert, "Vaccines do not cause autism, but I don't know that they don't." The reason that such assertions are generally unassertible is that the second conjunct carries a strong implicature that the reason for saying

---

[4]Douven (2006).
[5]Lackey (2007).

that one doesn't know is that one lacks adequate evidence.[6] But this can be cancelled by an adequate explanation on the speaker's behalf:

> "I don't mean to convey that there isn't excellent evidence that vaccines do not cause autism; I only say that I don't know because a recent situation has led me not to believe it. But to be perfectly clear, you *should* believe that vaccines don't cause autism. In fact, I think that I should also believe it, but I just can't at the moment."

One important upshot is that a reasons-based norm like SRNA can adequately explain the intuitively unassertibility of Moorean propositions. And just as with the lottery paradox, what seems to explain the unassertibility of knowledge versions of Moorean propositions (*i.e.,* "*p*, but I don't know that *p*.") is that one lacks adequate epistemic *reasons.* Therefore, one need only invoke a reasons-based norm: KNA is once again overkill. The other important upshot is that SRNA adequately explains potential cases of warrantedly assertible Moorean propositions. Moreover, insofar as there are plausible cases of warrantedly assertible Moorean propositions, KNA simply gets these cases wrong. On the one hand, in Chapter 6 I argued that the Moorean data should drop out of the debate over norms of assertion; instead, the data should be viewed as an adequacy condition for proposed norms of assertion: any candidate norm must at least be able to explain the data. On the other hand, I take the case I have built to give us reason, on balance, to think that the Moorean data actually supports a reasons-based norm over KNA when we take all of the Moorean data into consideration.

---

[6]See, for example, Rysiew (2007).

So SRNA seems again to score both a practical and theoretical point from the Moorean data: it can explain the whole range of Moorean data, and what underpins KNA's explanation is really the reasons-based story.

In Chapter 7 I discussed the linguistic data involving various challenges to assertions such as, "How do you know?" KNA advocates argue that knowledge challenges are appropriate by default in (nearly) every context, and what explains this is that knowledge is the norm of assertion. I have offered at least two responses to this line of argument. First, it's not clear that challenges that contain an epistemic term are actually challenging whether a speaker's assertion meets the epistemic standard indicated by the term. That is, it's not clear that "Are you certain?" is really challenging whether a speaker's assertion that p rises to the epistemic status of certainty. Second, challenges such as "How do you know?" seem to elicit as wholly adequate responses a speaker providing her (epistemic) *reasons* for the assertion. Once again, the KNA explanation for *why* "How do you know?" challenges knowledge is at the level of reasons: the inclusion of "how" tells us why. Notice that if the challenge were, "Do you know that?", then the speaker could respond with, "Yes." But if the hearer is unsatisfied, she's forced to challenge with, "How do you know?", to which the speaker will respond by merely offering her reasons. Justification, reasons, or evidence is doing all the heavy lifting in the KNA explanation of the data.

I have further argued that a reasons-based norm can explain the data at least as well as KNA can, including the general propriety of "How do you know?", which might otherwise seem an ace in the hole for the KNA advocate. In most conversational circumstances, we

expect speakers to assert only if they know, although this isn't actually the norm of assertion. This is why "How do you know?" is an appropriate challenge in (nearly) every context. But KNA advocates have argued that knowledge- and weaker-challenges really do challenge a speaker's *knowledge*: even if I'm correct that "How do you know?" is best interpreted as directly challenging a speaker's (epistemic) reasons for a proposition, since such reasons are a necessary condition for knowledge (*qua* justification), such challenges challenge knowledge. In a sense this is undeniable: challenging a necessary condition of x thereby challenges x. But it's only indirect, at best.

I have argued that based on the principle of wholly adequate responses, challenges to assertions require, at most, a speaker to articulate her (epistemic) reasons for her assertion. Even stronger challenges such as "Are you certain?" only require a speaker to give her reasons, albeit very good reasons. So if all sorts of challenges, especially the knowledge and stronger challenges most often invoked by KNA (and CNA) advocates seem to *directly* challenge (epistemic) reasons, and a reasons-based norm can adequately accommodate the data, then this seems to best support a norm like SRNA rather than KNA. Although KNA can conceivably explain the data, it does so at the cost of piggybacking on the SRNA explanation: KNA is once again overkill. So just as with the lottery and Moorean data, SRNA seems at least to enjoy a theoretical advantage over KNA, if not also a practical one.

9.2 Summing Up and Moving Forward

The overarching purpose of this dissertation was to advance the view that assertion is governed by a central *reasons-based* epistemic norm, which I have called the Supportive Reasons Norm (SRNA). I have offered the Knowledge Norm (KNA) as the principal foil for my view, since it is arguably the most influential view advanced in the recent literature. There are three primary sources of support offered for KNA: the lottery paradox, Moore's paradox, and challenges to assertion. In each case, though, I've argued that SRNA explains the linguistic data at least as well, if not better, than KNA. And, moreover, the underlying details of the KNA explanation for the unassertibility of some propositions seems to piggy-back on the reasons-based explanation. In other words, even when the KNA seems to explain and predict intuitions or linguistic data, its explanations are derivative on something like the SRNA. Therefore the KNA explanation is superfluous.

This is a source of theoretical support for SRNA over KNA: insofar as two views can equally well explain a set of data, then *ceteris paribus* the more parsimonious view should be preferred. But I also think that SRNA enjoys better practical support than KNA in that SRNA is able to account for some cases that KNA seems to simply get wrong. For example, KNA incorrectly predicts that *all* Moorean sentences are never warrantedly assertible and, moreover, that all false assertions are never warrantedly assertible. I've offered three cases of the latter: Stella (from Lackey's creationist teacher example), Jim the physics teacher, and Bill the taxi driver.

Since SRNA is not the only reasons-based norm on offer, it's important to distinguish my

view from others'. In Chapter 8 I distinguished my view from Lackey's RTBNA and Douven's RCR. Neither of the latter views adequately accounts for cases in which a speaker warrantedly asserts something she knows to be false (such as Jim the physics teacher). Moreover, neither RTBNA nor RCR integrates important social and pragmatic features of the context of assertion *into* the norm, something that SRNA does with condition (ii). While neither view precludes accommodating these features independently, I have argued that we should prefer an account that illuminates and unifies normative aspects of assertion, epistemic and otherwise.

There are a number of important questions left unanswered in my work, though. The project of understanding assertion and its norms has wider scope than merely understanding a single linguistic practice. For example, there's an emerging debate concerning the central epistemic norm of practical reasoning. On one view, the norm for practical reasoning and the norm for assertion are identical. Hawthorne & Stanley (2008) have proposed that KNA is that norm. However, I suspect that something like SRNA will work for practical reasoning as well. On such a view, one may use something as a premise in practical reasoning only if there are supportive reasons for the premise for the context.[7] For example, in deciding to p, one must have adequate reasons to believe that p will maximize one's expected utility for the context of decision.[8] In fact, there's reason to think that SRNA is better positioned by noticing that the social and pragmatic aspects of making decisions are likely different from those present in assertion. So it's best to have a norm that respects these differences, rather

---

[7]My future work will build on McKinnon (2011), where I criticize the use of lottery examples in the argument for the knowledge norm of practical reasoning.

[8]The context, for example, may suggest a satisficing rather than a maximizing strategy.

than glossing over them as KNA does. However, this is a direction for further inquiry.

References

Adler, J. (2002). *Belief's Own Ethics.* Cambridge, MA: MIT Press.

Alston, W. (1988). An internalist externalism. *Synthese*, 74(3), pp. 265 – 283.

Alston, W. (2005). Perception and representation. *Philosophy and Phenomenological Research*, 70(2), pp. 253 – 289.

Audi, R. (2006). Testimony, credulity, and veracity. In J. Lackey & E. Sosa (Eds.), *The Epistemology of Testimony* (pp. 25 – 49). Oxford: Oxford University Press.

Austin, J. (1946). Other minds. *Proceedings of the Aristotelian Society*, Supp. 20, 148–187.

Austin, J. (1962). *How to Do Things with Words: The William James Lectures delivered at Harvard University in 1955.* Oxford: Clarendon Press.

Bach, K. & Harnish, R. (1979). *Linguistic Communication and Speech Acts.* Cambridge, MA: MIT Press.

Benton, M. (2011). Two more for the knowledge account of assertion. *Analysis*, 71(4), 684–687.

Black, M. (1952). Definition, presupposition, and assertion. *Philosophical Review*, 61(4), 532–550.

Boghossian, P. (2008). Epistemic rules. *Journal of Philosophy*, 105(9), pp. 472 – 500.

Brandom, R. (1983). Asserting. *Noûs*, 17(4), pp. 637–650.

Brandom, R. (1998). *Making it Explicit: Reasoning, Representing, and Discursive Commitment.* Cambridge, MA: Harvard University Press.

Brown, J. (2008). The knowledge norm for assertion. *Philosophical Issues*, 18, 89 – 103.

Cappelen, H. & Lepore, E. (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism.* Malden, MA: Blackwell.

Carter, A. & Gordon, E. (2011). Norms of assertion: The quantity and quality of epistemic support. *Philosophia*, Forthcoming.

de Almeida, C. (2001). What moore's paradox is about. *Philosophy and Phenomenological Research*, 62(1), 489–523.

de Almeida, C. (2007). Moorean absurdity: An epistemological analysis. In J. N. Williams & M. Green (Eds.), *Moore's Paradox: New Essays on Belief, Rationality, and the First Person*. Oxford: Oxford University Press.

de Almeida, C. (2011). Epistemic closure, skepticism and defeasibility. *Synthese*, Forthcoming.

DeRose, K. (1991). Epistemic possibilities. *Philosophical Review*, 100(4), pp. 581 – 605.

DeRose, K. (1992). Contextualism and knowledge attribution. *Philosophy and Phenomenological Research*, 52, 913–29.

DeRose, K. (1996). Relevant alternatives and the content of knowledge attributions. *Philosophy and Phenomenological Research*, 56(1), 193–197.

DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review*, 111, 167–203.

Dodd, D. (2007). Why williamson should be a sceptic. *Philosophical Quarterly*, 57(229), 635–649.

Dougherty, T. & Rysiew, P. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78(1), pp. 123 – 132.

Douven, I. (2006). Assertion, knowledge, and rational credibility. *Philosophical Review*, 115(4), 449–485.

Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49, pp. 1 – 22.

Dummett, M. (1981). *Frege: Philosophy of Language*. London: Duckworth, 2 edition.

Fantl, J. & McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Fogelin, R. & Sinnott-Armstrong, W. (1997). *Understanding Arguments: An Introduction to Informal Logic*. Belmont, CA: Wadsworth, 5 edition.

Frankfurt, H. (1988). *The Importance of What We Care About: Philosophical Essays*. Cambridge: Cambridge University Press.

Gluer, K. & Pagin, P. (1999). Rules of meaning and practical reasoning. *Synthese*, 117(2), pp. 207 – 227.

Gluer, K. & Wikforss, A. (2010). The normativity of content and meaning. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. Stanford, CA: The Metaphysics Research Lab.

Goldberg, S. (2006). Reductionism and the distinctiveness of testimonial knowledge. In J. Lackey & E. Sosa (Eds.), *The Epistemology of Testimony.* Oxford: Oxford University Press.

Goldberg, S. (2007). *Anti-Individualism: Mind and Language, Knowledge and Justification.* Cambridge: Cambridge University Press.

Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20), pp. 771 – 791.

Goldman, A. (1986). *Epistemology and Cognition.* Cambridge, MA: Harvard University Press.

Grice, P. (1989). *Studies in the Way of Words.* Cambridge, MA: Harvard University Press.

Grover, D. (1992). *A Prosentential Theory of Truth.* Princeton, NJ: Princeton University Press.

Harman, G. (1973). *Thought.* Princeton: Princeton University Press.

Hawthorne, J. (2004). *Knowledge and Lotteries.* Oxford: Oxford University Press.

Hawthorne, J. & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.

Hendricks, V. (2006). *Mainstream and Formal Epistemology.* Cambridge: Cambridge University Press.

Hilger, M. & Taylor, I. (2007). *The Poker Mindset: Essential Attitudes for Poker Success.* Dimat Enterprises.

Hill, C. & Schechter, J. (2007). Hawthorne's lottery puzzle and the nature of belief. *Philosophical Issues*, 17, 102–122.

Horwich, P. (1998). *Truth.* Oxford, UK: Clarendon Press, 2nd edition.

Huemer, M. (2007). Moore's paradox and the norms of belief. In S. Nuccetelli & G. Seay (Eds.), *Themes from G.E. Moore* (pp. 142–157). New York: Oxford University Press.

Jary, M. (2010). *Assertion.* Houndmills, Basingstoke, Hampshire: UK: Palgrave Macmillan.

Kant, I. (1897). On the supposed right to lie from benevolent motives. `http://oll.libertyfund.org`. Last Accessed, January 25th, 2012.

Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes From Kaplan* (pp. 481–563). Oxford: Oxford University Press.

Kindred, D. (1999). Against all odds – hole-in-one records. `http://findarticles.com/p/articles/mi_m0HFI/is_1_50/ai_5430606`. Last Accessed, December 27, 2010.

Kvanvig, J. (2009). Assertion, knowledge, and lotteries. In P. Greenough & D. Pritchard (Eds.), *Williamson on Knowledge* (pp. 140–160). Oxford: Oxford University Press.

Kvanvig, J. (2011). Norms of assertion. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 233–250). Oxford: Oxford University Press.

Kyburg, H. (1961). *Probability and the Logic of Rational Belief.* Middletown, CT: Wesleyan University Press.

Lackey, J. (1999). Testimonial knowledge and transmission. *The Philosophical Quarterly*, 49, pp. 471 – 490.

Lackey, J. (2007). Norms of assertion. *Noûs*, 41(4), 594–626.

Lackey, J. (2008). What luck is not. *Australasian Journal of Philosophy*, 86(2), 255–267.

Lackey, J. (2010). Assertion and isolated second-hand knowledge. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 251–276). Oxford: Oxford University Press.

Lackey, J. & Sosa, E. (2006). *The Epistemology of Testimony.* Oxford: Oxford University Press.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.

Littlejohn, C. (2010). Moore's paradox and epistemic norms. *Australasian Journal of Philosophy*, 88(1), 79–100.

Maitra, I. (2011). Assertion, norms, and games. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 277 – 296). Oxford: Oxford University Press.

Mayr, E. (1991). *One Long Argument: Charles Darwin and the Genesis of Modern Evolutionary Thought.* Cambridge, MA: Harvard University Press.

McKinnon, R. (2011). Lotteries, knowledge, and practical reasoning. *Logos and Episteme: An International Journal of Epistemology*, 2(2), pp. 225 – 231.

McKinnon, R. (2012). How do you know that 'how do you know?' challenges a speaker's knowledge? *Pacific Philosophical Quarterly.* Forthcoming.

Moore, G. (1962). *Commonplace Book: 1919 - 1953.* London: Allen & Unwin.

Moore, G. E. (1942). *The Philosophy of G. E. Moore.* New York: Tudor. Paul Schilpp, ed.

Moran, R. (2006). Getting told and being believed. In J. Lackey & E. Sosa (Eds.), *The Epistemology of Testimony*. Oxford: Oxford University Press.

Plantinga, A. (1993a). *Warrant and Proper Function*. Oxford: Oxford University Press.

Plantinga, A. (1993b). *Warrant: The Current Debate*. Oxford: Oxford University Press.

Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.

Pryor, J. (2004). What's wrong with moore's argument. *Philosophical Issues*, 14(Epistemology), 349–378.

Quine, W. (1986). *Philosophy of Logic*. Cambridge, MA: Harvard University Press.

Rescorla, M. (2009). Assertion and its constitutive norms. *Philosophy and Phenomenological Research*, 79(1), 98–130.

Rysiew, P. (2007). Speaking of knowledge. *Noûs*, 41, 627–662.

Schaffer, J. (2008). Knowledge in the image of assertion. *Philosophical Issues*, 18, pp. 1–19.

Searle, J. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge, UK: Cambridge University Press.

Searle, J. (1979). *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge, UK: Cambridge University Press.

Shieber, J. (2009). Epistemological contextualism and the knowledge account of assertion. *Philosophia*, 37, pp. 169–181.

Siebel, M. (2003). Ilocutionary acts and attitude expression. *Linguistics and Philosophy*, 26(3), pp. 351 – 366.

Sklansky, D. (2007). *The Theory of Poker*. Henderson, NV: Two Plus Two Publishing, 4 edition.

Smith, M. (2010). What else justification could be. *Noûs*, 44(1), pp. 10 – 31.

Sorensen, R. (1988). *Blindspots*. Oxford: Oxford University Press.

Sorensen, R. (2007). Bald-faced lies! lying without the intent to deceive. *Pacific Philosophical Quarterly*, 88(2), 251–264.

Sosa, E. (2011). *Knowing Full Well*. Oxford: Oxford University Press.

Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18, 33–55.

Stine, G. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, pp. 249 – 261.

Stone, J. (2007). Contextualism and warranted assertion. *Pacific Philosophical Quarterly*, 88, 92–113.

Turri, J. (2010a). Does perceiving entail knowing? *Theoria*, 76(3), 197–206.

Turri, J. (2010b). Epistemic invariantism and speech act contextualism. *Philosophical Review*, 119(1), 77–95.

Turri, J. (2010c). Prompting challenges. *Analysis*, 70(3), 1–6.

Turri, J. (2011). The express knowledge account of assertion. *Australasian Journal of Philosophy*, 89(1), 37–45.

Unger, P. (1975). *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.

Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubting: Contemporary Perspectives in Skepticism*. Dordrecht: Kluwer Academic Publishers.

Vogel, J. (1999). The new relevant alternatives theory. *Philosophical Perspectives*, 13(13), 155–180.

Weiner, M. (2005). Must we know what we say? *Philosophical Review*, 114(2), 227–251.

Whewell, W. (1840). *The Philosophy of Inductive Sciences*. London: J. and J.J. Deighton.

Williams, J. (1994). Moorean absurdity and the intentional 'structure' of assertion. *Analysis*, 54(3), 160–166.

Williams, J. N. & Green, M., Eds. (2007). *Moore's Paradox: New Essays on Belief, Rationality, and the First Person*. Oxford: Oxford University Press.

Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105, 489–523.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

Wilson, E. (1998). *Consilience: The Unity of Knowledge*. New York: A. A. Knopf Inc.

Wright, C. (1992). *Truth and Objectivity*. Cambridge, MA: Harvard University Press.