

Assertion and its Constitutive Norms

Michael Rescorla

Abstract: Alston, Searle, and Williamson advocate *the restrictive model of assertion*, according to which certain constitutive assertoric norms *restrict* which propositions one may assert. Sellars and Brandom advocate *the dialectical model of assertion*, which treats assertion as constituted by its role in the game of giving and asking for reasons. Sellars and Brandom develop a restrictive version of the dialectical model. I explore a *non-restrictive* version of the dialectical model. On such a view, constitutive assertoric norms constrain how one must react if an interlocutor challenges one's assertion, but they do *not* constrain what one should assert in the first place. I argue that the non-restrictive dialectical perspective can accommodate various linguistic phenomena commonly taken to support the restrictive model.¹

§1. Constitutive assertoric norms

Assertion invites normative assessment from diverse perspectives. We may condemn Iago's lies to Othello as immoral, simultaneously commending how effectively they promote Iago's goals or how courteously Iago expresses himself. In rendering these judgments, we deploy evaluative standards rooted in morality, instrumental rationality, and propriety. The evaluative standards exhibit no special connection to language. We could readily apply them to Iago's non-linguistic actions. Assertion falls under general norms lacking any peculiarly linguistic character.

Does assertion also fall under peculiarly linguistic norms? Norms that issue, not from general principles of morality, rationality, or social interaction, but from special features of linguistic intercourse? Many philosophers answer affirmatively, arguing that assertion falls under *constitutive norms*. On this view, assertion is analogous to a game, in that it intrinsically involves certain rules. Someone who exhibits no sensitivity to the rules of a game does not grasp what it is to play the game, and someone who exhibits no sensitivity to assertion's constitutive norms does not grasp what it is to assert a proposition. Proponents of this view include Alston (2000), Brandom (1994), Douven (2006), Searle (1969, 2001), Sellars (1963), and Williamson (2000).

Williamson holds that assertion falls under a constitutive norm with the structure: One should assert p only if $C(p)$. Williamson's candidate for ' $C(p)$ ' is 'one knows p ,' yielding

The Knowledge Norm: One should assert only propositions that one knows.

Countless other norms with the same structure are possible, including:

The Honesty Norm: One should assert only propositions that one believes.

The Truth Norm: One should assert only true propositions.

The Warrant Norm: One should assert only propositions that one believes with 'sufficient' warrant.

I will call such norms *restrictive*, since they restrict which propositions one may assert. *The restrictive model of assertion* holds that assertion falls under constitutive restrictive norms.

Alston, Brandom, Searle, and Sellars endorse the restrictive model, as do most philosophers who embrace constitutive assertoric norms.

One can also imagine *non-restrictive* assertoric norms. I focus on norms with the structure: If one asserts p , then, if $D(p)$, one must $\phi(p)$. Theories based entirely upon norms with this structure reject any notion of 'assertibility' or 'appropriate assertion.' They characterize

assertion solely in terms of how it alters the speaker's normative standing, not in terms of which assertions are permissible.

The most promising candidates for non-restrictive assertoric norms reflect what I will call *the dialectical model of assertion*, which regards assertion as essentially a move within 'the game of giving and asking for reasons.' The intuitive idea here is that, by asserting a proposition, I commit myself to defending the proposition when faced with challenges and counter-arguments. I can cancel the dialectical commitment by retracting my assertion, but until then the commitment stands. Although the basic idea is ancient, recent discussion of it derives mainly from the writings of Hamblin (1971), Sellars (1963), and Toulmin (1958). However we elaborate this idea, the resulting account will likely feature non-restrictive norms that describe how to react if another speaker contests an asserted proposition.

One can develop the dialectical model in either a restrictive or non-restrictive direction. On the restrictive version, constitutive assertoric norms codify how to answer challenges and counter-arguments, and they *also* constrain which propositions one can initially assert. Brandom (1994), Sellars (1963), and Watson (2004) advocate views along these lines. I will explore a non-restrictive version of the dialectical model. On this view, constitutive assertoric norms constrain how I must react if someone challenges my assertion, but they do *not* constrain what I should assert in the first place. Morality, propriety, or rationality might require that I assert only propositions with some desired feature. No constitutive norm of assertion imposes any such requirement. For instance, a lie violates no constitutive assertoric norm, although the liar incurs dialectical commitments that she may or may not be able to discharge.

A non-restrictive version of the dialectical model may appear somewhat implausible. I am aware of only a single recent commentator, MacFarlane (2003, 2005), who advocates

anything resembling it. I will urge that some popular arguments for the restrictive model are seriously flawed. I will furthermore argue that the non-restrictive dialectical perspective can accommodate various linguistic phenomena commonly taken to support the restrictive model.

§2. Norms and practices

In this section, I develop a rough conception of how constitutive assertoric norms differ from non-constitutive assertoric norms. I hope that the conception will be fairly uncontroversial.

Advocates of constitutive assertoric norms typically emphasize the analogy with games, as in the following passage by Williamson: “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” (p. 238). The analogy with games seems helpful, so we should explore it in detail. How exactly is assertion more analogous to a game than to an activity such as jumping?

The contrast between playing a game and jumping illustrates a more general contrast between what I will call *practices* and *mere activities*. Besides games, sample practices include: dances, like the waltz or the tango; religious ceremonies, like weddings; fraternity initiation rites; performing Beethoven’s Fifth Symphony. Every practice is associated with “internal” standards of normative assessment codified by norms dictating how to execute the practice correctly. The norms for performing Beethoven’s Fifth Symphony demand that musicians play the notes of Beethoven’s score at roughly the indicated tempo. The norms for dancing the waltz demand that dancers move in time with the music, selecting their movements from a fixed repertoire of dance steps. Agents who violate a practice’s norms do not implement the practice correctly, but they

may still implement the practice. A devious gambler might repeatedly cheat during a poker game. He plays poker incorrectly, but he plays poker nonetheless.

The connection with internal evaluative standards differentiates practices from mere activities, like jumping, bathing, or holding hands. A mere activity falls under general norms, like those of morality, propriety, and rationality, but it does not fall under specialized norms that govern correct execution of the activity. There is no correct or incorrect way to bathe. There is no correct or incorrect way to hold hands.

The contrast between practices and mere activities suggests the following formulation: a norm is *constitutive of a practice* iff one must obey the norm to engage correctly in the practice.

In some weak sense, *any* activity involves constitutive requirements: requirements one must satisfy to count as engaging in the activity. One must immerse parts of one's body in some liquid to count as 'bathing.' Two people must achieve physical contact between their hands to count as 'holding hands.' But only practices enshrine constitutive standards that govern *correct* execution of the activity. 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. We might describe this contrast by saying that all activities involve *constitutive requirements*, which one must satisfy to engage in the activity, while practices also involve *constitutive norms*, which one must satisfy to engage in the practice correctly.

There are basically two ways one might analyze constitutive assertoric norms within the framework developed thus far. Most straightforwardly, we might treat assertion as a self-contained practice governed by constitutive norms. On this view, constitutive assertoric norms

describe how to engage correctly in the practice of assertion, where that practice can be characterized in isolation from other speech acts. Alston, Searle, and Williamson adopt this approach. A somewhat less straightforward approach treats assertion as a move within a *larger* practice, much as a home run is a move within the larger practice of baseball. On this view, constitutive assertoric norms describe how to engage correctly in some segment of the larger practice. Sellars and Brandom adopt this approach, insisting that assertion is individuated by its role within the game of giving and asking for reasons, which I will often refer to as ‘reasoned discourse.’ If Sellars and Brandom are correct, then assertion can be fully characterized only through its normative relations to other speech acts, such as questioning and challenging.

Although I will side with Sellars and Brandom in this dispute, our terminology should not take sides. I offer the following characterization. A norm is a *norm of assertion* iff my assertoric performances are relevant to whether I violate it. A norm is *constitutive of assertion* iff it is a norm of assertion *and* it is constitutive of a practice to which assertion is intrinsically connected. This formulation is rough, but it is general enough to encompass Alston, Searle, and Williamson (since the practice might simply be assertion itself) and Sellars and Brandom (since the practice might be the game of giving and asking for reasons). Our formulation reflects the intuitive idea that assertion has an intrinsically normative dimension. We can characterize assertion’s fundamental nature only by delineating specialized norms that govern its correct execution. I think this conception guides many philosophers, either more or less explicitly.

On the other hand, many philosophers repudiate specialized assertoric norms. Davidson is a particularly emphatic opponent (1984, pp. 265-280). In my terminology, Davidson regards conversation as a *mere activity*, rather than as a *practice* incorporating its own internal standards of normative assessment. For Davidson, there is no significant distinction between ‘correctly’ or

‘incorrectly’ executing speech acts. One either executes them or one does not. Such an approach precludes anything like constitutive norms of speech acts. It countenances only whatever normativity results from subsuming conversation under extra-linguistic norms. On Davidson’s approach, assertion is not intrinsically normative, although, like any action, it falls under norms of morality, propriety, rationality, etc.

My discussion highlights the elementary but crucial distinction between a norm’s *content* and its *normative status*. When discussing ‘norms of assertion,’ it is never enough merely to specify a norm’s content, i.e. what the norm demands. One must also specify what normative force it exerts. *In some sense*, it is clearly a ‘norm of assertion’ that one ‘should’ speak honestly. The question is whether the Honesty Norm enjoys a *specifically linguistic* grounding beyond what it receives from morality, legality, propriety, rationality, and other extra-linguistic factors. Is the Honesty Norm *constitutive* of assertion? Or does it merely result from subsuming assertion under more general normative patterns?

§3. The dialectical model of assertion

According to the dialectical model, assertion involves a commitment to defend the asserted proposition if challenged. The model develops this idea, in part, by positing constitutive norms that regulate how one should respond when other speakers challenge one’s assertions.

Does ‘defending’ my assertion require arguments that *actually* support my position, or that I regard as doing so, that my interlocutor regards as doing so, or something else? Can my argument employ any non-question-begging premises, or only true premises, or only premises my interlocutor accepts, or something else? How we answer such questions determines which norms we embrace. For instance, selecting the first answer to each question suggests

The Defense Norm: When challenged to defend an asserted proposition, one must either provide a cogent, non-circular argument for the proposition or else retract it, where a ‘cogent’ argument is one whose premises provide rational support for its conclusion. Often, proponents of the dialectical model regard the Defense Norm as too strong, preferring instead something closer to

The Default-Challenge Norm: When faced with a *legitimate* challenge to defend an asserted proposition, one must either provide a cogent, non-circular argument for the proposition or else retract it.

For example, a challenge to the assertion “I have hands” might count as legitimate only if the challenger sketches a compelling skeptical scenario in which I do not have hands. Leite (2005) presents a theory of reasoned discourse based on something like the Default-Challenge Norm. He does not endorse the further claim that this norm is constitutive of assertion. Brandom endorses the further claim. In (Rescorla, 2009b), I defend a weakened variant of the Defense Norm against Leite and Brandom. We may ignore these internecine skirmishes here.

The dialectical model also describes how one should respond to counter-arguments against an asserted proposition. Again, many variations are possible. Most accounts feature something like the following norm:

The Retraction Norm: When faced with a counter-argument against an asserted proposition, one must rebut the counter-argument or else retract the proposition.

Note that we can accept the Retraction Norm even if we reject anything resembling the Defense or Default-Challenge Norm, a combination of views suggested by (MacFarlane, 2005, p. 334).

If I seek to discharge a dialectical commitment, as opposed to canceling it, then I must provide an argument. How do I provide an argument? *By asserting additional propositions* as the

argument's premises. Thus, assertion plays a dual role in reasoned discourse: through it, one both undertakes *and* discharges dialectical commitments. As Brandom puts it:

[A]ssertings... are in the fundamental case what reasons are asked for, and what giving a reason always consists in. The kind of commitment that a claim of the assertional sort is an expression of is something that can stand in need of (and so be liable to the demand for) a reason; and it is something that can be offered as a reason (1994, p. 167).

Assertion's dual role might appear to generate a vicious regress, since discharging dialectical commitments embroils one in further commitments. Proponents of the dialectical model evince various reactions to this regress. Brandom tries to halt the regress through the Default-Challenge Norm, which allows us to say that certain assertions require defense only given special effort by one's interlocutor. In contrast, (Rescorla, 2009a) argues that the regress is not worrisome: either speaker and interlocutor agree upon mutually acceptable relevant premises, in which case the regress halts, or speaker and interlocutor do *not* agree upon mutually acceptable premises, in which case the speaker leaves certain dialectical commitments undischarged.² Again, we may ignore these complications here.

To undercut the dialectical model, some critics emphasize that serious dialectical interaction is a relatively rare phenomenon. We often assert propositions without engaging one another in reasoned discourse. Yet this undeniable fact is consistent with the dialectical model, which holds that asserting a proposition is performing an action that occupies a certain normative role within reasoned discourse. Such a view can allow that assertion sometimes occurs outside reasoned discourse. Non-dialectical assertoric performances may even *statistically outnumber* performances within reasoned discourse. The model claims only that non-dialectical assertoric performances are *explanatorily* derivative from core performances within reasoned discourse.

Part of what makes a given linguistic performance an assertion, as opposed to some other speech act, are the dialectical commitments the speaker assumes, commitments that become operative once the speaker engages fellow conversationalists in reasoned discourse. In this sense, assertion is ‘intrinsically connected’ to reasoned discourse.

Not all assertions are moves within reasoned discourse, but all assertions are *potential* moves within reasoned discourse. Whenever I assert some proposition, other speakers may challenge me to defend it or even provide counter-arguments against it. And then I face a choice. I can extrude myself from reasoned discourse (“I’m not going to argue with you”). I can engage in reasoned discourse *incorrectly*, trying to defend my assertion but doing so inadequately. Or I can engage in reasoned discourse *correctly*, vindicating my assertion or else retracting it. Thus, the norms of reasoned discourse describe a normative infrastructure that lurks beneath the surface of quotidian conversation, perpetually threatening to emerge. According to the dialectical model, assertion is individuated by its place within this normative infrastructure and hence is possible only against the background provided by it.

But aren’t there are many assertoric performances through which I undertake no dialectical commitment whatsoever? Suppose I engage someone in casual small-talk at a cocktail party. I might assert a proposition to initiate a new topic or promote a diverting conversation. Yet, one might urge, it hardly seems appropriate to greet such an assertion with challenges or counter-arguments. So the dialectical model must be fundamentally mistaken in tying assertoric force so closely to dialectical commitment.

In response, note that there are many elliptical gambits through which I can introduce a new topic or promote an entertaining conversation. Instead of asserting p , I might say, “I’m curious what you think about p .” I thereby *put forward* a proposition p without undertaking any

dialectical commitment. If my interlocutor responds by arguing against p , I might simply reply, “Very interesting points.” It would be outrageous for my interlocutor to then demand, “Well, aren’t you going to respond to my arguments? Or, at least, retract what you said?”. That demand seems far more legitimate if I initially asserted p . Although the demand may be obnoxious or uncouth, there is a clear sense in which it is invited by an assertion that p but not by more elliptical ways of putting forward p . This contrast suggests the following analysis: assertion always involves some element of dialectical commitment, but general norms of social propriety may render it more or less appropriate to press the speaker to discharge those commitments.

Most proponents of the dialectical model combine it with the restrictive model. Brandom treats the Warrant and Default-Challenge Norms as constitutive of assertion. Watson (2004) develops a similar account, supplemented with a restrictive norm along these lines:

The Defensibility Norm: One should assert only propositions one can defend with cogent, non-circular arguments.

We can imagine numerous variants, such as:

The Known Defensibility Norm: One should not assert propositions one knows one cannot defend with cogent, non-circular arguments.

In what follows, I will investigate the prospects for a non-restrictive version of the dialectical model. I will examine various popular arguments for the restrictive model, urging that any uncontroversial phenomena adduced by them can be accommodated just as well by the non-restrictive dialectical perspective. For additional development of the non-restrictive dialectical perspective, see (Rescorla, 2007).

§4. Does the dialectical model entail the restrictive model?

One might suspect that the dialectical model collapses into the restrictive model, and hence that the non-restrictive dialectical perspective is incoherent. If assertion involves a commitment to defend what one says, then shouldn't one avoid asserting propositions one cannot defend, or at least propositions one knows one cannot defend? More explicitly, one might reason:

- (A) (i) By asserting a proposition, one commits oneself to defending it.
- (ii) One should not undertake commitments one knows one cannot discharge.
- (iii) Hence, one should not assert propositions one knows one cannot defend.

Argument (A) purports to derive a restrictive norm from the basic intuition (i) underlying the dialectical model. If 'defend' means 'defend with cogent, non-circular arguments,' then the derived norm is the Known Defensibility Norm.

In evaluating (A), a central consideration is the normative force of the 'should' in (ii) and (iii). Undoubtedly, one should not undertake commitments one knows one cannot keep, *if one wants to keep one's commitments*. But this establishes only that

If one wants to keep one's commitments, one should not assert propositions one knows one cannot defend,

which is a hypothetical norm about how to achieve a certain goal (keeping one's commitments), not a constitutive norm of assertion. One might instead construe (ii) as a moral norm. Under this construal, (A) is valid only if 'should' in (iii) likewise expresses moral obligation, in which case (iii) is moral norm rather than a norm describing how to execute assertion correctly.

Another idea would be to treat (ii) as somehow constitutive of commitment *in general*, in which case it seems more plausible that (i) and (ii) jointly yield a constitutive norm of assertion. But, as I will now argue, this suggestion also faces problems.

Consider Rawls on *promising*. Rawls holds that promising falls under a constitutive norm that “reads roughly as follows: if one says the words ‘I promise to do X’ in the appropriate circumstances, one is to do X, unless certain excusing conditions obtain” (1971, p. 345). Rawls explains our obligation to keep promises as resulting from interaction between this norm and the Principle of Fairness, a general moral prohibition against free-riding on just practices. He does *not* mention a constitutive norm against promising to ϕ when one has no intention of ϕ -ing or even when one knows one cannot ϕ . A natural interpretation, then, is that Rawls does not think insincere promises violate any constitutive norm of promising. Violation occurs when one fails to keep the promise, but not necessarily before. By promising to ϕ when I have no intention of ϕ -ing, I may be *morally* culpable. But moral transgressions are not necessarily violations of constitutive promissory norms. For instance, suppose I promise to ϕ while feeling quite certain I will not be able to ϕ ; shortly after the promise, it becomes both possible and advantageous for me to ϕ , so I ϕ . On the proposed view, I need violate no constitutive norm of promising, although I undoubtedly violate a moral norm.

Alston (2000) and Searle (1969) argue that promising falls under a restrictive constitutive norm that one promise to ϕ only if one intends to ϕ . Other philosophers, such as Thomson (1990, p. 303), question whether promising involves constitutive norms at all. I take no stand regarding which account is correct. My point is just that there is nothing *incoherent* about the conception I have attributed to Rawls. One can imagine a practice along these lines, even if that practice is not ours, and even if there are good reasons for preferring our own. The proposed practice provides a coherent model of commitment as regulated by norms requiring one to fulfill one’s commitments but no norms restricting which commitments one undertakes. It thereby casts doubt upon the

thesis that (ii) is constitutive of commitment in general, and hence upon (A)'s attempt to derive a constitutive restrictive norm from the dialectical model.

There is *something* defective about performing action ψ within practice P if I know doing so will eventually force me to engage in P incorrectly. But that does not entail that ψ itself, rather than some action subsequent to ψ , violates P 's constitutive norms. The following principle is plausible: engaging in a practice constitutively requires at least pretending that one seeks to engage in it correctly. Thus, depending on the details of the case, my ψ -ing is defective in one of two ways. Either I do not really seek to engage in P correctly, in which case my behavior is disingenuous, since I misleadingly pretend that I seek to engage in P correctly. Or else I seek to engage in P correctly, in which case it is pragmatically incoherent to ψ , since I know doing so will eventually force me to violate P 's constitutive norms. In neither case need we posit an additional constitutive norm, a norm violated by ψ , beyond those constitutive norms ψ will eventually induce me to violate.

The non-restrictive dialectical perspective may not yet seem plausible. But it *does* seem coherent. In particular, constitutive features of commitment do not induce a collapse of the dialectical model into the restrictive model.

§5. Representing oneself as believing/knowing that p

I can speak insincerely. It may even be common knowledge between myself and my interlocutor that I am speaking insincerely (Rumfitt, 1995). But I cannot assert a proposition while abandoning all pretensions towards speaking sincerely. If I say "Frank is an excellent student" while blatantly smirking and winking, then I do not assert that Frank is an excellent student, except perhaps in some attenuated sense that does not interest us here. My body

language subverts the display of candor that is a necessary concomitant to genuinely, non-ironically asserting a proposition. Assertion essentially involves at least a pretense of sincerity.

It is tempting to suppose that the essential link between assertion and apparent sincerity supports the restrictive model. If assertion essentially involves a *display* of sincerity, then surely one should assert a proposition only when the display is genuine.

This reasoning is fallacious. It slides without any argument from uncontroversial claims about assertion's *constitutive requirements* to highly controversial claims about assertion's *constitutive norms*. A constitutive requirement of assertion is that one adopt a show of conviction, just as a constitutive requirement of bathing is that one immerse part of one's body in liquid. In neither case does it follow that the activity involves *constitutive norms*: norms specifying how to execute the activity correctly. Someone who abandons the pretense of sincerity does not engage incorrectly in assertion. She does not assert a proposition at all. Of course, given that assertion involves a display of sincerity, and given the moral principle that one should avoid deceiving others, it follows that one *should* speak sincerely. But that still does not show that any norm is constitutive of assertion. It only illustrates that assertion falls under general norms that apply to many other actions.

Few philosophers would embrace the fallacious reasoning in the simplistic form just criticized. But similar mistakes plague much more sophisticated expositions.

Unger (1975, p. 253) advocates the following thesis: "If someone asserts, states, or declares that something is so, then it follows that he represents himself as *knowing* that it is so." Let us call this *Unger's Thesis*. According to DeRose (2002, p. 180), Unger's Thesis entails that the Knowledge Norm is constitutive of assertion: "For our purposes, these are just two sides of

the same coin: If one represents oneself as knowing that p by asserting p , then, to avoid falsely representing oneself, one should follow the rule of asserting only what one knows.”

I am not convinced that Unger’s Thesis is true. But the key point here is that Unger’s Thesis does not entail that the Knowledge Norm is constitutive of assertion. At best, Unger’s Thesis identifies a *constitutive requirement* of assertion: to assert p , one must represent oneself as knowing p . That assertion embodies this constitutive requirement does not entail that there are correct or incorrect ways of asserting a proposition. Only a confusion between constitutive requirements and constitutive norms, or between constitutive and non-constitutive norms, could make the inference from Unger’s thesis to the restrictive model seem obvious and unproblematic. Davidson, who endorses a weakened version of Unger’s Thesis that replaces *knowledge* with *belief*, puts the point succinctly: “What we do in making an assertion is represent ourselves as believing what we say, and so we may be morally at fault if it turns out that we don’t believe what we asserted. But moral error is not linguistic error” (2005, p. 123).

Hence, there is no deductive entailment from Unger’s Thesis, or similar doctrines, to the restrictive model of assertion. But one might urge that some kind of *abductive* entailment holds. In the previous section, we mentioned the following extremely plausible principle: engaging in a practice constitutively requires at least pretending that one seeks to engage in it correctly. This principle, combined with the Honesty Norm, entails that assertion constitutively requires at least a pretense of sincerity. Similar derivations are possible for Unger’s Thesis, or for the weakened version of it espoused by Davidson (Williamson, 2000, p. 252). So the restrictive model can explain assertion’s constitutive requirements.

One problem with the proposed abductive argument is that one might simply deny that assertion’s constitutive requirements need explaining. Davidson urges that his weakened version

of Unger's Thesis follows directly from the conceptual analysis of assertion (1984, p. 270). It is not clear that the proposed abductive argument exerts any force against Davidson's position.

However, the abductive argument poses a challenge to the non-restrictive dialectical perspective. The argument shows that the restrictive model easily accommodates relatively uncontroversial features of assertion. It is not clear that a non-restrictive dialectical account can explain those same features. Apparently, I can undertake a commitment to defend some proposition without any pretense of believing it. For instance, I might promise a friend that I will defend some proposition against all challenges and counter-arguments, even while emphasizing that I do not believe the proposition. How, then, can the non-restrictive dialectical perspective explain why assertion requires apparent sincerity?

In §3, I noted that assertion occupies a dual role in reasoned discourse: through it, one both undertakes *and* discharges dialectical commitments. So far, I have emphasized the first of these two roles. As I will now argue, focusing on the second role helps defuse the present objection to non-restrictive dialectical theories. Schematically: while apparent sincerity may not be necessary for *undertaking* dialectical commitment, it is necessary for *discharging* dialectical commitments; so assertion's role in discharging dialectical commitment generates a constitutive requirement of apparent sincerity.

I begin by exploring more closely the nature of "dialectical commitment." In developing the dialectical model, we should not depict reasoned discourse as a rhetorical exercise utterly detached from the participants' beliefs. By asserting a proposition, I make it *my* position, not merely a position I agree to defend for the sake of argument. Assertoric commitment is not just a commitment to providing arguments or to rebutting counter-arguments. As I will put it, assertion involves a commitment to *advocate* for one's position, i.e. *to present it as worthy of belief*.

Advocacy is not included in what the Defense and Default-Challenge Norms demand. Specifically, it requires me to present my arguments *as* providing compelling reasons for believing the asserted proposition. For instance, it is not enough merely to say, “Here are some arguments that *p*, but I don’t myself find them plausible.” If I issue this disclaimer, I cease to advocate for the asserted proposition. (A lawyer might tell the jury, “Here are some arguments for my client’s innocence. I don’t find them convincing, but you should make up your own mind.” The lawyer does not advocate for her client, no matter how impressive her arguments.)

How should we capture ‘advocacy’ through precise norms? I employ a fairly minimalist formulation: advocacy involves presenting the premises of one’s argument *as* providing reason for believing the asserted proposition. So construed, I think that advocacy underlies several familiar versions of the dialectical model, including (Brandom, 1994). It is compatible with various dialectical norms. For instance, we could supplement the Defense Norm with the proviso that, when providing cogent arguments, one must advocate for one’s position. We could add an analogous proviso to a norm demanding arguments that *one’s interlocutor* finds cogent, or that *the speaker* finds cogent.

According to the dialectical model, assertion is individuated by its role in reasoned discourse. Part of this role consists in the fact that speakers discharge dialectical commitments by asserting propositions. So assertion, by its very nature, is a means for discharging dialectical commitments. Clearly, I need not intend each assertoric performance as contributing to the discharge of a dialectical commitment. I might assert *p* on a whim, without envisaging a broader argumentative goal. Nevertheless, there is a clear sense in which even this whimsical assertion contributes to the future putative discharge of dialectical commitments. Metaphorically: the assertion requires no supplementation for me to cite *p* as a premise in future arguments. I need

merely re-affirm my prior assertion. This element of redundancy differentiates assertion from speech acts like *conjecturing p* or *putting p forward*. If I conjecture *p* or put it forward, then I cannot subsequently present it as affording reason for belief without fundamentally transforming my standing towards it. If I assert *p*, no such transformation is required. By asserting *p*, I have already presented it a potential premise in future arguments I may or may not offer.

As Brandom puts it, then, “[u]ttering a sentence with assertional force or significance is putting it forward *as* a potential reason... Assertions are essentially *fit* to be reasons” (1994, p. 168). Part of what makes my speech act an assertion of *p*, rather than a conjecture or a mere putting forward of *p*, is that I present *p* as providing reason for believing other propositions. There may not be any specific proposition that I advance *p* as affording reason for believing. But I present it as a reason to be invoked, when relevant, in future advocacy.

We may now offer the following argument that assertion constitutively requires a pretense of sincerity:

- (B)**
- (i) By asserting a proposition, one presents it as providing reason for believing other propositions.
 - (ii) One presents a proposition as providing reason for believing other propositions only if one maintains some minimal pretense of believing it.
 - (iii) Hence, one asserts a proposition only if one maintains some minimal pretense of believing it.

(i) and (ii) entail the desired conclusion (iii). The previous four paragraphs argued that (i) seems plausible once we accept a suitable version of the dialectical model. (ii) seems plausible whatever our conception of assertion. Suppose I utter *p* while indicating, perhaps through winks or other body language, that I do not believe *p*. Then how can I present *p* as providing reason for

believing any other proposition q ? I can advance it as a reason why *some other person* (including my interlocutor) might believe q , but not as a reason why *one should* believe q . Lacking a minimal show of credence in p , I cannot depict it as having any reason-giving force.

Revisiting an earlier example, suppose I promise a friend that I will vigorously defend p , even while emphasizing to my friend that I do not believe p . Although I thereby undertake a kind of dialectical commitment, I do not present p as affording reason for believing any other proposition. Thus, my speech act does not contribute to the future putative discharge of dialectical commitments. So it is not an assertion.

In motivating (i), we invoked structural features of reasoned discourse that any plausible version of the non-restrictive dialectical perspective should accept. Moreover, (i) articulates a constitutive requirement, not a constitutive norm, so there is no obvious entailment from (i) to claims about assertion's constitutive norms. Hence, (B) neither presupposes nor entails that assertion falls under constitutive restrictive norms. In this way, our argument shows how a non-restrictive dialectical account can explain why assertion requires apparent sincerity.

To buttress (B), imagine a dialectical practice that is “devil’s advocate all the way down.” In this hypothetical practice, I can advance a proposition p , thereby incurring a commitment to defend it. I defend p by advancing additional propositions as premises in an argument for p . My interlocutor may in turn challenge those premises. However, neither my initial utterance of p nor my subsequent defense of it need involve any pretense of sincerity. I merely advance and defend propositions for the sake of argument, as when participants in our own dialectical practice employ the following locutions: “Well, I’m not saying that I believe this myself, but some would argue as follows,” “Proponents of p would respond to your counter-argument by saying q ,” etc.

Could the hypothetical practice exist autonomously from normal assertoric practice? For instance, must not participants in any stable, coherent dialectical practice have available a speech act through which they at least pretend to believe claims about word-meanings or justificatory relations between propositions? Leaving such questions aside, it seems clear that the hypothetical practice differs dramatically from our own dialectical activity. The key difference is that, even if participants in the hypothetical practice offer cogent arguments, they do not advance those arguments as providing reasons for belief. In this sense, the practice hardly counts as a “game of giving and asking for *reasons*.” Participants in the practice may provide arguments, but they do not advocate for propositions. They undertake a fundamentally different kind of dialectical commitment than that which we undertake through assertion. Apparently, then, any dialectical practice remotely like ours requires a speech act through which one undertakes dialectical commitments *and* throughout the performance of which one maintains some pretense of belief.

It is likely that reasoning similar to (B) can explain Davidson’s weakened version of Unger’s Thesis, according to which someone who asserts p represents herself as believing p . Specifically, we might replace (ii) with the following premise: one can present a proposition as providing reason for believing other propositions only if one represents oneself as believing it. What about the original version of Unger’s Thesis, based upon knowledge rather than belief? We might invoke the much stronger premise (ii’): one can present a proposition as providing reason for believing other propositions only if one represents oneself as knowing it. Although I do not myself find (ii’) particularly congenial, it strikes me as no less plausible than Unger’s Thesis. Unger himself (1975, pp. 206-214) advocates a position somewhat reminiscent of (ii’): if X ’s reason for believing something is that p , then X knows that p . Williamson (2000, pp. 184-208) advances a kindred doctrine: only known propositions are evidence for other propositions. Thus,

while the topic obviously requires further discussion, it hardly seems clear that Unger's Thesis poses insuperable difficulties for the non-restrictive dialectical perspective.³

§6. Committing oneself to the truth of what one asserts

Peirce remarks that assertion involves an essentially normative ingredient: "This ingredient, the assuming of responsibility, which is so prominent in solemn assertion, must be present in every genuine assertion" (1934, p. 386). One can formulate Peirce's insight in various ways: a speaker *commits himself* to the proposition asserted; he *takes responsibility* for its truth; he *endorses* the proposition; he *stakes himself* to the truth of the proposition; and so forth.

Many philosophers suggest that these Peircean formulations support the restrictive model. For instance, Searle urges that "asserting commits the speaker to the truth of the proposition asserted" (2001, p. 147). He infers that the Truth Norm is constitutive of assertion. As he puts it, "[s]uch principles as 'you ought to tell the truth,' 'you ought not to lie,' or 'you ought to be consistent in your assertion' are *internal* to the notion of assertion. You do not need any external moral principle to have the relevant commitments. The commitment to truth is built into the structure of the intentionality of the assertion" (2001, 181).

An obvious objection to Searle's argument is that it rests upon the somewhat obscure notion 'committing oneself to the truth of a proposition.' As MacFarlane (2005, p. 334) notes while developing this objection, the most straightforward notion of commitment is commitment to *do* something. I commit myself to attending a conference, to paying my Visa bill, and so on. But an assertion that *p* is not, in general, a commitment to make it the case that *p*. In what sense, then, do I 'commit' myself to the truth of a proposition that I assert?

One might propose that the relevant commitment is to *believing* the proposition, or perhaps to *judging* that it is true. But what is it to undertake such a commitment? I can promise to accumulate evidence and to survey it judiciously, but I can hardly promise what conclusions I will draw. I can commit myself to *self-inducing* the belief that *p*, perhaps by swallowing a pill with appropriate belief-altering properties. Obviously, this commitment is not an element of normal assertion. On the other hand, if the putative commitment is to *its already being the case that I believe or judge p*, then this revives the question of how to understand a commitment that is not commitment to some course of action.

Another proposal is that, as Searle puts it, “[i]n making an assertion we take *responsibility* for truth, sincerity, and evidence... These responsibilities are met only if the world is such that the utterance is true, the speaker is sincere, and the speaker has evidence for the assertion” (2001, p. 176). The relevant notion of “taking responsibility” is still fairly obscure. I can take responsibility for *doing* something, such as watering my neighbor’s plants while he vacations. I can take responsibility for having caused some proposition to be true. But what is it to take responsibility for a proposition being true when I played no role in its truth? I can take responsibility for compensating my interlocutor if what I assert turns out not to be true. Yet this obviously has no bearing upon typical assertoric performances. Another idea, explored by Alston (2000, p. 54-63), is to construe ‘*x* takes responsibility for its being the case that *p*’ in terms of the censure, blame, or reproach that *x* would or should encounter if it turns out that *x* did not believe *p*. I will discuss these negative reactive attitudes in §7. Alston eventually rejects construals based upon reactive attitudes, adopting instead the following analysis: “*x* takes responsibility for its being the case that *p* in uttering *s* iff, in uttering *s*, *x* subjects his utterance to a rule that implies it is permissible for *x* to utter *s* only if *p*” (p. 60). He argues that this rule-based analysis is more

explanatorily fundamental than one based upon negative reactive attitudes, since it explains the source of those attitudes and also their pattern of distribution across various cases. However, the rule-based analysis embodies something quite close to the restrictive model. Thus, as Alston would doubtless acknowledge, it cannot subserve an *argument* for the restrictive model.

In ordinary conversation, we frequently adduce a speaker's 'conversational commitments.' We say that a speaker is committed to some proposition, or to some proposition being true. We criticize a speaker who cannot defend his conversational commitments when faced with challenges or counter-arguments. Thus, I think that Searle evokes fundamental intuitions about assertion. But the non-restrictive dialectical perspective arguably accommodates these Peircean intuitions just as well as the restrictive model. We can gloss 'commitment to the truth of a proposition' as 'commitment to *defending* (or *advocating for*) the truth of a proposition.' On this construal, a speaker who is 'committed to a proposition' or 'committed to the truth of a proposition' is committed to defending the proposition within reasoned discourse. Crucially, a commitment to defending the truth of what one says is not the same as a commitment to speaking the truth in the first place.

It might seem that the proposed dialectical analysis provide a natural entrée for the restrictive model. Commitment to defending the truth of what one says may not be literally the same as commitment to speaking the truth. The former commitment may not even entail the latter. But they seem to go together as a natural package. Why would I incur the former commitment if I don't also incur the latter?

The adversarial legal system provides an instructive, albeit imperfect, analogy. As instantiated within the United States, this practice is governed by a complex set of constitutive norms, some codified and others implicit in the common law tradition. Two lawyers advocate for

competing sides, with the verdict decided by a neutral third party. The lawyers are committed to defending their positions as effectively as possible. Blatant failure to do so constitutes malpractice and merits sanction. Yet Applbaum (2000) argues that advocacy within this system falls under no global Honesty Norm. As Applbaum puts it, “over the course of representing a client, a lawyer forms reasonable beliefs... about the various factual propositions at issue in the trial, and... zealous advocacy often requires her to attempt to persuade the jury to believe the opposite” (p. 106). The relevant ‘reasonable beliefs’ sometimes rise to the level of knowledge. Legal norms do not allow lawyers to say just anything. Complex rules of evidence forbid certain especially egregious deceptions. But no generalized Honesty Norm prevails. Thus, there is a clear sense in which legal advocates are not, in general, committed to speaking the truth.

From the non-restrictive dialectical perspective, assertion exhibits many similarities with the adversarial legal system. By asserting a proposition, I commit myself to advocating for it, just as a lawyer who takes a case commits herself to advocating for her client’s position. There is no constitutive commitment to advocate only for positions one believes, only a commitment to advocate for whatever position one chooses. In neither case, then, does commitment to defending the truth of what one says entrain commitment to speaking the truth in first place. Of course, there are many disanalogies between assertion and the adversarial system: legal practice involves explicitly codified norms directly sustained by the threat of punitive sanctions; there are supposedly neutral adjudicators distinct from the advocates; part of practice’s rationale is the need to preserve individual rights; and so on. But the non-restrictive dialectical approach holds that these disanalogies conceal a more fundamental homology in normative structure.

This picture of assertion strikes many philosophers as absurd. It is admittedly a bit extreme. So far, though, we have found no cogent argument against it.

§7. Intuitions of speaker-defectiveness

Philosophers often motivate the restrictive model by adducing pre-theoretic intuitions of *defectiveness*. We may divide the relevant intuitions into two categories: intuitions that *the speaker* is somehow blameworthy; and intuitions that, whether or not the speaker is culpable, her *assertion itself* is defective.⁴ I will describe these intuitions respectively as *speaker-defect intuitions* and *assertion-defect intuitions*. A liar might elicit speaker-defect intuitions. A speaker who asserts a false proposition she has excellent reason to believe might elicit assertion-defect intuitions. I discuss speaker-defect intuitions in this section and assertion-defect intuitions in §8.

If we find that a speaker lied, we may blame or censure her. The restrictive model can explain these negative reactive attitudes as registering that the speaker deliberately violated some constitutive assertoric norm. But do the negative reactive attitudes reflect anything special about assertion? Or do they reflect general norms that encompass both assertion and many other actions? A useful technique here is to elicit comparative judgments regarding analogous cases of assertoric and non-assertoric communication. Non-assertoric counterparts function somewhat like controls in a scientific experiment.

Consider the following examples, several of which derive from two famous cases discussed by Kant: the “murderer at the door,” and the man who packs his luggage to create the misleading impression he is embarking on a trip.

1(a): John is helping an innocent victim evade a killer who also is John’s neighbor. The victim is hiding in John’s car. To fool the killer, who does not know that John is harboring the victim but who is standing far down the street, John goes to his car with a suitcase and shouts to the killer, “I’ll be out of town for a few days,” then drives away. He does not intend to leave town.

1(b): Similar to 1(a), except that this time John says nothing to the neighbor. He packs a heavy suitcase and carries it to his car. He waves to his neighbor, and he points silently to the suitcase while loading it in the car. He does this to create the same misleading impression as in 1(a).

2(a): While preparing to drive to a marathon poker game at a nearby location, John shouts to his nosy but harmless neighbor far down the street, “I’ll be out of town for a few days.” He lies because he does not want his neighbor to know why he will be gone.

2(b): Stands to 2(a) as 1(b) stands to 1(a).

3(a): Iago tells Othello that Desdemona betrayed him with Cassio.

3(b): Iago shows Othello doctored photographs of Desdemona betraying him with Cassio.⁵

In each pair of examples, an agent conveys roughly the same proposition. But he conveys it *assertorically* in (a) and *non-assertorically* in (b). If the Honesty, Truth, Warrant, or Knowledge Norm is constitutive of assertion, then (a) exhibits a specifically linguistic defect beyond whatever defects it shares with (b). Unlike (b), (a) violates the constitutive norms of assertion.

Does the deceiver in (a) exhibit a qualitatively different type of defect than in (b)? I do not think so. Our relevant intuitions within each pair of examples strike me as commensurate. We view 1(a-b) as commendable, 2(a-b) as involving at most a fairly minor infraction, and 3(a-b) as involving a serious infraction. Assertoric communication is quite different than non-assertoric communication, so we might cite many disanalogies between each (a) scenario and the corresponding (b) scenario. But it is hard to see how these disanalogies secure any *intuitive, pre-theoretic* respect in which (a) elicits a fundamentally different censorious reaction than (b). Thus, in explaining (a)’s blameworthiness, or lack thereof, there is no need to invoke specialized assertoric norms. We can invoke general evaluative standards that apply both to assertion and to

many other actions. When assessing the deceiver for praise or blame, the fact that he employed *assertion* as his mode of deception seems negligibly significant.

An imposing philosophical tradition, famously represented by Aquinas and Kant, insists that direct lies are somehow worse than other forms of deception.⁶ Bernard Williams, who accuses Aquinas and Kant of “fetishizing assertion,” even so concedes that under certain circumstances “there is something peculiarly odious or insulting about a lie as contrasted with other forms of deceit” (2002, p. 118). The most popular illustration is *conversational implicature*. As Augustine observed, there are many contexts in which misleading conversational implicatures seem far less objectionable than outright lies. On this basis, Adler (1997) argues that assertion, unlike conversational implicature, falls under a constitutive norm of truthfulness. As another example, imagine a variant of scenario 2(b) in which John carries his luggage to his car but pretends not to notice his neighbor observing him. This scenario, which is similar to how Kant himself intended the luggage example, seems markedly less “odious and insulting” than a direct lie. At the very least, such examples suggest that a lie’s defects outstrip its role in disseminating propositions the speaker believes false. A lie is not bad simply because it deceives or because it is intended to deceive.

Nevertheless, we should not conclude that lies are *in general* more deplorable than other forms of deceit. That conclusion seems plausible only if we focus on a narrow range of unrepresentative examples. For instance, the deceived neighbor in Kant’s version of the luggage example does not know that John seeks to influence her beliefs. The example therefore involves nothing resembling common knowledge of communicative intentions and expectations. Once we incorporate this common knowledge, as in 2(b), John’s deceit becomes far more objectionable.

Now, John not only dupes his neighbor, he also abuses her trust. Given this abuse of trust, 2(b) seems no less “odious and insulting” than 2(a).

Misleading conversational implicature is likewise an unrepresentative example. Conversational implicature inherently lends itself to craftiness. It is a way of conveying propositions without committing oneself to them. If I conversationally implicate *p* rather than asserting *p*, it is natural for my interlocutors to wonder *why*. Why do I beat around the bush, rather than just come out and say what I mean? In many contexts, then, the decision to implicate *p* rather than assert *p* signals *caveat auditor*. As Adler himself observes, “the choice of a deception rather than a lie can itself convey a message of moral effort and struggle” (1997, p. 452). Adler’s observation helps explain why misleading conversational implicatures often seem less offensive than outright lies. Unlike the liar, a speaker who misleads through conversational implicature provides an implicit warning, through his choice of communicative strategy, that we should parse his message with extreme care. If his deception does not seem *so* bad, that is partly because his mode of deception already broadcasts his guile.

In contrast, the modes of deception employed in scenarios 1-3(b) raise no red flags. We may construe each scenario so that it seems quite natural for the deceiver to communicate his message non-assertorically: in 1-2(b), because John’s neighbor is quite far away; and in 3(b), because the photographs speak for themselves, requiring no verbal supplementation. Once construed in this way, 2-3(b) evince intuitive defects comparable to 2-3(a). Thus, excessive focus upon conversational implicature, at the expense of more suitable examples, encourages the mistaken impression that assertoric deceit is peculiarly repugnant.

My goal is not to analyze why lying is wrong or how it relates to other forms of deception. As the ethical literature on mendacity demonstrates, a complete analysis would

require sophisticated casuistical maneuvers applied to a wide range of cases. A good account might invoke some of the following familiar ideas: deception often yields bad consequences for the deceived; the deceiver abuses a bond of trust with the deceived; the deceiver acts on a maxim that is not universalizable; the deceiver treats the deceived as a means, not as an end; the deceiver violates our transpersonal rational interest in the pursuit of truth; the deceiver undermines our trust in one another and thereby diminishes our ability to communicate; and so on. My point is that, despite these complexities, we lack any antecedent reason to expect that a good account will stigmatize assertoric deception as uniquely blameworthy when compared to suitable varieties of non-assertoric deceit. A special pejorative stance towards assertoric deception may be a *consequence* of a satisfying theory, but it is not a *pre-theoretic constraint*.

I turn now to cases where the speaker believes the asserted proposition but lacks “sufficient” grounds for asserting it. Such cases often elicit the intuition that the speaker, while not as culpable as a liar, behaves irresponsibly. The restrictive model can explain this intuition by saying that the speaker exerts insufficient effort to obey constitutive assertoric norms.

I begin with an example made famous by Williamson. The lottery results will be announced in a few minutes, and my friend is waiting to hear whether her ticket won. She purchased the ticket on a lark, so she is not very invested in the outcome. I say, “Your ticket did not win.” Call the proposition asserted “the lottery proposition.” My grounds for believing the lottery proposition are purely probabilistic. I do not have any inside information about the lottery. According to Williamson, I do not know the lottery proposition, even though I am justified in believing it. Williamson also claims that “you will still be entitled to feel some resentment when you later discover the merely probabilistic grounds for my assertion” (p. 246). Williamson argues that we should explain this intuitive reaction through the Knowledge Norm.

Given the details of the lottery example, it is difficult to construct a clean non-assertoric counterpart akin to 1-3(b). But we can study some related cases:

4(a): John's company is planning a weekend retreat in the country. The company is about to announce the result of a lottery, conducted last night. The lottery's winner, if any, will be excused from the retreat. Everyone else must attend. John is leaving for the office to learn the lottery's result. If his ticket won, he is determined to stay in town. If not, he is determined to go on the retreat. He knows his ticket has very low odds of winning (let them be as low as you please). He brings luggage with him so that he can leave directly for the retreat if necessary. As he enters his car, he shouts to his neighbor, "I'll be out of town for a few days."

4(b): Similar to 4(a), except that John says nothing. Instead, he packs a heavy suitcase, waves at his neighbor, and points silently at the suitcase while loading it into his car. He does this to communicate that he will be out of town for a few days.

It seems fairly clear that John's basis for believing he will leave town is probabilistic. By Williamson's lights, John does not know he will leave town. Suppose that, in both 4(a) and 4(b), John's neighbor eventually learns about the lottery and its significance. Williamson's account predicts that the assertoric case merits more resentment than the non-assertoric one. I think it far from clear that pre-theoretic intuition validates this prediction. As long as John's odds of winning the lottery are low enough, I am not sure that 4(a) elicits any resentment at all. Even if it does, I do not think it seems any *worse* than 4(b). Both cases strike us as, at worst, mildly irresponsible but hardly grounds for strident criticism.

We can imagine a series of variants upon 4(a) and 4(b), in which we alter the nature and strength of John's warrant for believing that he will be out of town for a few days. In each pair of cases, I conjecture that John elicits no more resentment in (a) than in (b).

Generally, the literature sparked by Williamson's discussion has debated *which* restrictive assertoric norm best explains the lottery example, not *whether* we should explain the example through a restrictive norm. Weiner (2005) argues that we should explain the lottery example through the Truth Norm and Gricean implicature, while Kvanvig (2009) favors an account based on the Warrant Norm. Such treatments, like Williamson's, typically cite speaker-defect intuitions without paying even cursory attention to non-assertoric counterparts. In my view, this is somewhat like trying to explain why property *X* causes cancer before checking that subjects with *X* acquire cancer more frequently than subjects lacking *X*. Before documenting intuitive discrepancies with comparable non-assertoric cases, we should not assume that negative reactions elicited by assertoric performances reflect anything special about assertion.

The parity between 1-4(a) and 1-4(b) suggests that speaker-defect intuitions provide little support for the restrictive model. Any resentment evoked by 1-4(a) reflects not specialized assertoric norms but general norms governing belief-transmission through both linguistic and non-linguistic means. Admittedly, I have not formulated those general norms, nor is it obvious how to do so. But there is no evident reason why the patterns of blame and praise surrounding 1-4(a) require a special framework of restrictive constitutive norms. On the contrary, that framework predicts disparities that find no echo in pre-theoretic intuition.

In response to my argument, proponents of the restrictive model might propose that 1-4(b) themselves fall under constitutive restrictive norms. On this proposal, speaker-defect intuitions register no disparity between (a) and (b) because they *both* violate constitutive norms.

The challenge facing this proposal is to explain why 1-4(b) involve a practice with the features delineated in §2: internal evaluative standards codified by constitutive norms. For instance, in 1-2(b) and 4(b), John communicates that he will be out of town for a few days by

exploiting his neighbor's mastery of folk psychology, not mastery of some pre-existing practice. Given John's conduct, it is common knowledge between John and his neighbor that he wants the neighbor to notice his suitcase. The most natural explanation for why John is loading a heavy suitcase into his car is that he is embarking on a trip that will last several days. So it is reasonable for the neighbor to infer that John is trying to instill the belief that he is embarking on a trip that will last several days. Since the neighbor believes from prior interactions with John that he is honest and reliable, the neighbor concludes that John is embarking on a trip that will last several days. This explanation does not depend in any obvious way upon constitutive norms.

§8. Intuitions of assertion-defectiveness

Assertion-defect intuitions record that the *assertion itself* is defective, even if the speaker is not culpable. As Alston puts it, "the utterance is *wrong*, *incorrect*, or *out of order*, even if the epistemological conditions for blameworthiness are not satisfied... [I]t was the wrong thing to say, whatever the correct judgment as to the speaker's fault... [T]here was something fundamentally amiss with the utterance" (2000, p. 57, 60). The point here is not just that some assertion is, say, false. That is something all sides can acknowledge. The point is that describing an assertion as false constitutes a *criticism* of it. By asserting a falsehood, the speaker fails to meet some relevant evaluative standard, no matter how hard she tried to reach it. As Alston puts it, "we must distinguish the perception of an out-of-order sentence utterance from the judgment that the assertion, or what is asserted, is false" (2000, p. 267). Similarly, one might claim that an assertion seems 'amiss' if the speaker does not believe what she asserts, or does not believe it with sufficient warrant, or does not know it. The restrictive model can explain these intuitions as registering that the speaker, perhaps unwittingly, violated some constitutive assertoric norm.

A basic difficulty here is that assertion-defect intuitions are much less robust than speaker-defect intuitions. Suppose that, mistakenly believing it is raining, I assert that it is raining. We can all agree that my assertion is false. But is it clear that this verdict constitutes a criticism? Do we really have a ‘perception’ that the utterance is ‘out-of-order,’ over and above our recognition that it is false? Matters become even less clear for 1(a), in which the assertion’s falsity and the speaker’s dishonesty might well seem the basis, not for criticism, but for unqualified acclaim. Of course, advocates of the Truth or Honesty Norms insist that 1(a) seems intuitively defective *qua* assertion, despite any countervailing moral or prudential merits. But the putative ‘intuition’ here seems perilously close to an intuition that the restrictive model is correct. Opponents of the restrictive model, such as Davidson, claim not to share such intuitions.

How can we decide between the restrictive and non-restrictive perspectives if we cannot agree upon the data to be explained? Do we face a Kuhnian stalemate in which rival theorists adduce divergent theory-laden intuitions? To some extent, perhaps. Before giving up, though, I will examine a few representative arguments advanced by proponents of the restrictive model. I will urge that those arguments assign assertion-defect intuitions a problematic role.

Consider the following passage from Searle (2001, p. 184):

Suppose I am not lying, but am genuinely mistaken. I sincerely said it is raining, but all the same it is not raining. In such a case there still is something wrong with my speech act, namely, it is false. But why is that wrong? After all, for every true proposition there is a false one. It is wrong because the aim of a statement is to be true, and this one fails, because it is false. When I make a statement I commit myself to its truth, and here my mistake makes me fail in my commitment.

This passage moves in a tight circle between theory and data, rendering the line between premise and conclusion difficult to discern. Apparently, Searle seeks to elicit assertoric-defect intuitions by invoking two ideas: first, assertion involves a commitment to the truth of what one says; second, the aim of assertion is truth. I critiqued the first idea in §6. The second idea seems too close to the restrictive model to provide a non-circular argument for it. Setting these criticisms aside, it seems clear that, rather than motivating the restrictive model by citing assertion-defect intuitions, Searle motivates the intuitions by citing putative normative features of assertion. But then Searle can hardly count the intuitions as pre-theoretic data upon which we should all agree.

Unlike Searle, Alston emphasizes that assertion-defect intuitions are elusive. His solution is to examine first language acquisition (2000, p. 267):

We seek to get across to the small child that he is not to say “A rabbit is outside” unless there is a rabbit outside... We try to get across to him when it is and when it is not all right to for him to utter a certain sentence... The details of first language acquisition are enormously complicated. But one thing that is crucially involved is the disposition on the part of teachers to spot cases in which the learner utters a sentence in violation of an I-rule and apply appropriate correction,

where ‘I-rule’ is Alston’s term for ‘constitutive norm of some speech act.’ Alston concludes that ordinary practice enshrines a ‘perception’ that false assertions are ‘out-of-order,’ a perception made manifest when we inculcate children into linguistic practice.

Alston observes that we encourage children to strive for true assertions. In fact, few parents teach their children that false assertions are never ‘all right,’ since lies are sometimes appropriate. Ignoring this quibble, Alston’s observation still does not support the restrictive model. As I argued in §5, the non-restrictive dialectical perspective can explain why assertion

requires a pretense of sincerity. Since asserting p requires a pretense that one believes p , asserting p is liable to induce others to believe p . Since moral norms require one, *ceteris paribus*, to avoid instilling false beliefs in others, they require one, *ceteribus paribus*, to strive to assert the truth. Hence, encouraging children to strive for true assertions might just be one chapter in their moral education. So Alston's observation does not suggest that a false assertion exhibits any defect beyond whatever moral, epistemic, cognitive, or rational defects may have led to it.

Alston also observes that, when we discover a child has asserted a falsehood, we 'apply appropriate correction.' Yet how do we apply this 'appropriate correction'? Prototypically, we confront the child with evidence that he asserted a false proposition and require him to concede, at least tacitly, that he was wrong. The non-restrictive dialectical perspective can explain this interaction as an instance of reasoned discourse, in which we try to induce the child to obey the Retraction Norm. The fact that we expect someone to retract an assertion does not suggest that the original assertion was defective, only that refusal to retract it would be defective.

We must distinguish two theses. The first thesis is that, when a speaker's interlocutors produce decisive counter-arguments against her assertion, then she must retract it. The second thesis is that the original assertion was itself somehow defective. Our intuitions strongly support the first thesis: if a speaker simply will not listen to reason ("That's my story and I'm sticking to it") then we find her linguistic conduct highly defective. But that suggests only something like the Retraction Norm. It does not follow that false assertions, or even assertions known by the speaker to be false, violate some constitutive assertoric norm.

Following Searle and Alston, I have focused upon cases in which a speaker unwittingly asserts a false proposition. What about cases where the speaker believes a true proposition but apparently lacks 'sufficient' grounds for asserting it? Consider again the lottery proposition.

Although Williamson recognizes the distinction between intuitions of speaker- and assertion-defectiveness (p. 261), he often blurs it, as in the following passage: “Intuitively, my grounds are quite inadequate for that outright unqualified assertion... You will still be entitled to feel some resentment when you later discover the merely probabilistic grounds for my assertion. I was representing myself to you as having an authority to make the flat-out assertion which in reality I lacked” (p. 246). The first and third sentences of this passage strike me as ambiguous between speaker- and assertion-defectiveness, while the second sentence evokes speaker-defectiveness. Taken as a whole, the passage confers spurious credibility upon assertion-defect intuitions by eliding the distinction with speaker-defect intuitions. This is no isolated expository anomaly. Discussions of the lottery example frequently either elide the distinction or else focus exclusively upon speaker-defect intuitions ((Weiner, 2005) illustrates the latter option).

Another worry centers again upon the crucial sentence: “You will still be entitled to feel some resentment when you later discover the merely probabilistic grounds for my assertion.” This sentence naturally invites us to consider an extended scenario in which I fail to defend the lottery proposition against challenges. It therefore elicits intuitions about how I should behave *after* asserting the lottery proposition, which are easily confused with intuitions about the assertion itself. For instance, suppose friend asks, “Why do you think my ticket didn’t win? What’s your evidence?”. I might respond by citing the low probability that the ticket won. But my friend could retort: “I know all about the probabilities. What I want from you is an argument that my ticket *did not* win, not an argument that it has a low chance of winning.” And to this I would lack any adequate reply. I would fail to vindicate my position. Williamson regards this failure as *evidence* of linguistic defectiveness. But one might instead regard the failure, assuming it is accompanied by a refusal to retract, as itself the sole linguistic defect. It is not

evidence for a further transgression of ‘exceeding one’s evidential authority,’ a transgression that already occurred with the original assertion. We must distinguish one undeniable intuition, that I should retract my assertion when I prove unable to defend it, from a further, much more debatable intuition that the original assertion was defective.

In general, it seems far from clear that we experience robust assertion-defect intuitions for the lottery example, or for other examples where the speaker asserts a proposition without ‘sufficient’ grounds. We may have an intuition that the speaker behaves irresponsibly, and we may have an intuition that the speaker must retract his assertion after failing to defend it. I argued in the previous section that we should explain the former intuition through general moral and social factors, not through constitutive assertoric norms. The latter intuition suggests only something resembling the Defense or Default-Challenge Norm. I submit that we can explain any non-controversial intuitions elicited by the lottery example, and other similar examples, through general moral and social norms combined with non-restrictive dialectical norms.

My analysis of the lottery example is consistent with almost any *epistemological* treatment. In particular, I take no stand regarding whether we *know* the lottery proposition or even whether we are *justified* in believing it. My analysis cites only an inability to provide an adequate *argument* supporting that proposition. As Alston (1989) and Audi (1993) emphasize, the ability to produce arguments for some proposition is quite distinct from the state of being justified in believing the proposition, let alone the state of *knowing* it.

Does my treatment of the lottery case presuppose that probabilistic arguments never discharge dialectical commitments? That presupposition would be tendentious. For instance, John might defend the assertion that he has some disease by citing the results of a blood test. The test has an extremely slight chance of false positives, so John’s defense appears probabilistic. Yet

the defense will strike many as perfectly legitimate. Why might we deem probabilistic argumentation acceptable in this scenario but not in the lottery example?

An adequate answer to this question would require a thorough analysis of probabilistic justification, a vexed topic in contemporary epistemology.⁷ Even lacking such an analysis, there is clearly something special about the lottery scenario that sets it apart from more routine cases like the blood test, as reflected by our unwillingness to say that one *knows* the lottery proposition but not that one can know one has a disease based on a blood test. Abductive considerations may help differentiate the two cases. The best explanation for why John's blood test is positive is that he has the disease; so a defense based on the blood test can appeal not just to 'bare' probabilities but to abduction. Abductive considerations are irrelevant to the lottery proposition, which is a prediction based upon good odds, not an explanation of anything.⁸ Thus, one might urge that probabilistic arguments are legitimate when they potentially admit abductive supplementation. Whether or not this maneuver generalizes to all other intuitively acceptable probabilistic arguments, it shows that I am not committed to any global denunciation of probabilistic argumentation in general. My analysis assumes only that, given the specifics of the lottery scenario, citing probabilities does not constitute an adequate defense of the lottery proposition.

Ultimately, I simply reject the putative assertion-defect intuitions that Alston, Searle, and Williamson seek to explain. However, our discussion suggests that assertion-defect intuitions are non-optimal candidates for pre-theoretic linguistic data. Although suggestive, they are theory-laden and labile. They are also easily confused with other quite different intuitions. This does not render them useless. An adequate theory of assertion must catalogue and account for them. But it must do so with greater circumspection than proponents of the restrictive model often exercise.

§9. The non-restrictive dialectical perspective

Evidently, it is more difficult than one might suppose to provide a compelling argument for the restrictive model. Although the literature contains many arguments besides those I have discussed, I believe that my discussion casts favorable light upon the non-restrictive dialectical perspective.⁹ That perspective is much more appealing than one would suspect based on how few contemporary philosophers advocate it. We should investigate it in more detail. At the very least, such investigations may bring into sharper relief the strengths and weaknesses of the restrictive model.

¹ I am greatly indebted to helpful feedback from Jonathan Adler, Anthony Corsentino, Kevin Falvey, Michael Glanzberg, Warren Goldfarb, Elizabeth Harman, Richard Heck, Louis-Philippe Hodgson, Charles Parsons, and James Pryor. I received many useful comments when I presented earlier versions of this paper at a UCSD Philosophy Colloquium, especially from Agustín Rayo and Wayne Martin, and at the Pacific APA 2005. I am grateful to Gary Ebbs, my commentator on the latter occasion, for his criticisms. Finally, I thank an anonymous referee for this journal for unusually generous and helpful comments.

² Does this approach entail that, if a speaker cannot satisfactorily defend p against a recalcitrant interlocutor, then the speaker is not justified in believing p ? No. My position concerns norms of assertion, not norms of belief. There are no easy inferences from the former to the latter. As Alston (1989) and Audi (1993) argue, we must sharply distinguish the ability to provide arguments for p from the state of being justified in believing p .

³ One complication is that “presenting x as y ” is an opaque context. Thus, even if we agree with Williamson that knowledge=evidence, we cannot immediately infer that one presents a proposition as evidence for other propositions only if one represents oneself as knowing it; similarly for any putative inference from Unger’s position to (ii’). Nevertheless, it may be that the arguments offered by Unger and Williamson, when suitably modified, support something sufficiently close to (ii’).

⁴ Roughly speaking, speaker- versus assertion-defect intuitions correspond respectively to what DeRose (2002) calls “secondary” versus “primary” violations of constitutive assertoric norms and what Alston (2000) calls shirking of one’s “subjective” versus “objective” obligations.

⁵ (Scanlon, 1990, p. 202-3) employs an analogous argumentative strategy, vis-à-vis promising rather than asserting. (Scanlon, 1999, p. 317-322) explains the defectiveness of lying in terms of generic moral considerations that also apply to non-linguistic communication. See also the discussion of promising in (Cavell, 1979, p. 298).

⁶ For an overview of this tradition, with citations, see (MacIntyre, 1995).

⁷ For discussion, with citations to the literature, see (Nelkin, 2000). Nelkin argues that bare probabilistic grounds are not sufficient for rational belief. She argues that beliefs based upon non-deductive grounds can nevertheless be rational when those grounds incorporate some abductive element.

⁸ Cf. (Nelkin, 2000, p. 404-408).

⁹ Alston (2000, p. 78), DeRose (2002), (Williamson, 2000, p. 253), and many others argue that Moore’s paradox supports the restrictive model. In contrast, Shoemaker (1996, p. 74-96) and Moran (2001) argue that Moore’s paradox is fundamentally a phenomenon at the level of *cognition*, rather than *language*. The locus of deviancy lies not with the *assertions* “ p , but I don’t know that p ” or “ p , but I don’t believe that p ,” but with the corresponding *beliefs*. As Shoemaker argues, we should first explain why there is something deviant about *believing* these propositions. Once we do so, it is unlikely that we will require substantial philosophical machinery to explain why there is something deviant about *asserting* them.

Another important observation is that one can appropriately answer any assertion by saying, “How do you know?”. Unger deploys this observation to motivate Unger’s Thesis, and Williamson deploys it to motivate the constitutive status of the Knowledge Norm. For critical discussion, see (Kvanvig, 2009).

Works Cited

- Adler, J. 1997. Lying, deceiving, or falsely implicating. *Journal of Philosophy* 94, 435-452.
- Alston, W. 1989. *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca: Cornell University Press.
- . 2000. *Illocutionary Acts and Sentence Meaning*. Ithaca: Cornell University Press.
- Appelbaum, A. 2000. *Ethics for Adversaries*. Princeton: Princeton University Press.
- Audi, R. 1993. *The Structure of Justification*. Cambridge: Cambridge University Press.
- Brandom, R. 1994. *Making it Explicit*. Cambridge: Harvard University Press.
- Cavell, S. 1979. *The Claim of Reason*. Oxford: Oxford University Press.
- Davidson, D. 1984. *Inquiries into Truth and Interpretation*. Cambridge: Cambridge University Press.
- . 2005. *Truth and Predication*. Cambridge: Harvard University Press.
- DeRose, K. 2002. Assertion, knowledge, and context. *Philosophical Review* 111, 167-203.
- Douven, I. 2006. Assertion, knowledge, and rational credibility. *Philosophical Review* 115: 449-485.
- Hamblin, C. L. 1971. Mathematical models of dialogue. *Theoria* 37, 130-155.
- Kvanvig, J. 2009. Assertion, knowledge, and lotteries. *Williamson on Knowledge*, ed. D. Pritchard and P. Greenough. Oxford: Oxford University Press.
- Leite, A. 2005. A localist solution to the regress of justification. *Australasian Journal of Philosophy* 83, 395-421.
- MacFarlane, J. 2003. Future contingents and relative truth. *The Philosophical Quarterly* 53, 250-65.
- . 2005. Making sense of relative truth. *Proceedings of the Aristotelian Society* 105, 321-39.
- MacIntyre, A. 1995. Truthfulness, lies, and moral philosophers. *The Tanner Lectures on Human Values, Volume XVI*, ed. G. B. Peterson. Salt Lake City: Utah University Press.
- Moran, R. 2001. *Authority and Estrangement*. Princeton: Princeton University Press.
- Nelkin, D. 2000. The lottery paradox, knowledge, and rationality. *Philosophical Review* 109, 373-409.
- Pierce, C. S. 1934. *Collected Papers*, vol. V. Cambridge: Harvard University Press.
- Rawls, J. 1955. Two concepts of rules. *Philosophical Review* 63, 3-32.
- . 1971. *A Theory of Justice*. Cambridge: Harvard University Press.
- Rescorla, M. 2007. A linguistic reason for truthfulness. In *Truth and Speech Acts*, ed. D. Greimann and G. Siegart. New York: Routledge.
- . 2009a. Epistemic and dialectical regress. *Australasian Journal of Philosophy* 87, 43-60.
- . 2009b. Shifting the Burden of Proof?. *The Philosophical Quarterly* 59, 86-109.
- Rumfitt, I. 1995. Truth-conditions and communication. *Mind* 104, 827-862.
- Scanlon, T. M. 1990. Promises and practices. *Philosophy and Public Affairs* 19, 199-226.
- . 1999. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Searle, J. 1969. *Speech Acts*. Cambridge: Cambridge University Press.
- . 2001. *Rationality in Action*. Cambridge: MIT Press, 2001.
- Sellars, W. 1963. Some reflections on language games. Rpt. in Sellars’s *Science, Perception, and Reality*. New York: Routledge and Kegan Paul.
- Shoemaker, S. 1996. *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.
- Thomson, J. J. 1990. *The Realm of Rights*. Cambridge: Harvard University Press.
- Toulmin, S. 1958. *The Uses of Argument*. Cambridge: Cambridge University Press.
- Unger, P. 1975. *Ignorance*. Oxford: Clarendon Press.
- Watson, G. 2004. Asserting and promising. *Philosophical Studies* 117, 55-77.
- Weiner, M. 2005. Must we know what we say? *Philosophical Review* 114, 227-251.
- Williams, B. 2002. *Truth and Truthfulness*. Princeton: Princeton University Press.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.