

Draft – Please do not cite without author’s permission.

Assertion, Norms, and Games

Ishani Maitra

October 11, 2009

§ 1 Introduction

In this paper, I focus on a package of views widely held by philosophers of language, and others. That package consists of the following three views.

- (1) Assertions are governed by either an alethic or an epistemic norm, i.e., a norm that specifies that it is appropriate to assert something only if what is asserted is true, or justifiably believed, or certain, or known.¹ There are substantive disagreements about which of these is *the* norm, as well as about how best to understand the concepts used in specifying the norms (e.g., truth, knowledge). Nevertheless, there is broad consensus that some such norm governs the speech act of assertion.²
- (2) One of the norms mentioned in (1) is more “intimately connected” to assertion than other norms regulating the speech act. (The phrase *intimately connected* is borrowed from Timothy Williamson (2000: 239).) The comparison class here includes norms such as humorousness (i.e., *Assert something only if it is funny*), tactfulness (i.e., *Assert something only if it is tactful*), sincerity, relevance, and clarity, to mention but a few. Each of the norms in this comparison class governs assertion in at least some contexts; some of them (e.g., relevance) may even govern assertion in every context. But, on the view being considered, there is a sense in which some alethic or epistemic norm bears an intimate connection to assertion that most (or all) of the other norms mentioned here don’t.³
- (3) The sense of intimate connection mentioned in (2) can be understood with the help of an analogy between language use and games. The analogy gets its start in the (now familiar) thought that language use is, in certain respects, like a game, and that assertion is a move within this larger game. Pursuing this analogy, it is suggested that the intimate

¹ Here, and throughout the paper, I will understand norms as standards for appropriate praise or criticism. To say that a norm governs assertion is to say that it is appropriate (in some sense) to praise assertions that satisfy that norm, and to criticize assertions that fail to satisfy it.

² The consensus isn’t universal. Some have recently defended norms that are neither alethic nor epistemic. For example, both Kent Bach (forthcoming) and Frank Hindriks (2007) have defended a belief norm for assertion, Igor Douven (2006) has defended a rational credibility norm, and Jennifer Lackey (2007) has defended a reasonable-to-believe norm.

Note that the phrase *epistemic norm* is sometimes used in a broad sense to include any competitor to what I will later call the *knowledge norm*. On this broad usage, all of the norms I’ve just mentioned, as well as alethic norms, are epistemic norms. I’m not going to be using *epistemic norm* in this very broad way.

³ Versions of this view are held by Michael Dummett (1978, 1993), Huw Price (1998, 2003), John Searle (1969, 1979), Jason Stanley (2007), Timothy Williamson (1996, 2000), and Crispin Wright (1992), among others.

connection between some norm(s) and assertion can be modeled upon the relationship between certain norms and games.⁴

In this paper, I will argue against the package consisting of the three views described above. More specifically, I intend to show that at least (3) should be rejected: even if some alethic or epistemic norm is intimately connected to assertion, that relationship is not mirrored by the relationship between norms and the games they govern.

Before proceeding to the criticisms, it will be useful to say something about what motivates (3) in the first place. There are, I think, two ways to motivate (3), corresponding to two ways of unpacking the analogy between language use and games.

First, for any game, what it is to play that game at all, rather than some other game (or no game), is constituted by certain norms. If a baseball player doesn't touch first base when rounding the bases, he may be playing baseball badly. If he doesn't go to first base after being awarded a walk, he isn't playing baseball at all. Some norms governing baseball, like the one that requires players to round the bases after hitting the ball, are such that it isn't possible to violate them *flagrantly* (in a sense to be explained in what follows) and still be playing the game in question. I shall say that such norms are *constitutive norms*.

Second, for any competitive game, some norms are connected to purposes that players have *qua* players of the game. The thought here is fairly simple. Competitive games have winning conditions. The purpose of a player, at least *qua* player, is to win. If he doesn't attempt to win at all, then he isn't playing the game. If a baseball player just stands between second and third base while ground balls roll by him, he isn't just playing the shortstop position badly, he isn't playing shortstop at all. So, when a player plays a competitive game, he must have winning among his purposes. (Some qualifications are necessary here, but I shall return to them later in the paper.) When a player has a particular purpose, he should do what he can to fulfill that purpose. This generates a host of norms. I shall say that norms that are derived from purposes that a player must have in order to count as playing the game are *purposive norms*.

Any (competitive) game is governed by a variety of norms. Among these, constitutive and purposive norms have a special place. First, they apply to players' actions in all contexts in which the game is played. Second, to know what it is to play a particular game (rather than some other game, or no game at all) is to know its constitutive and purposive norms. And third, every player must regard such norms as regulating their behavior when playing the game. Thus, we might say that the constitutive and purposive norms of a game are intimately connected to it in a way that other norms governing the game are not.

These observations suggest two ways of unpacking the analogy between language use and games. First, we might suppose that some norms governing a speech act are constitutive norms of the speech act, just as there are constitutive norms governing moves within a game. Second, we might suppose some norms governing a speech act are purposive norms of the speech act, just as there are purposive norms governing moves within games. If some norms are connected

⁴ This analogy is endorsed by Dummett, Searle, and Williamson, who agree on little else, though they unpack the analogy in different ways.

to assertion in either of these ways, that would be enough to bear out the idea that they are intimately connected to assertion in a way that other norms governing the speech act are not. Thus, we have (3).

One of the attractions of (3) is that it promises an answer to what I shall call *the hard question*: what kind of mistake does an asserter make when she violates a norm that is intimately connected to assertion? Putting the same question another way: from what sort of defect does an assertion that violates such a norm suffer? Suppose, for argument's sake, that an alethic norm (i.e., *Assert something only if it is true*) is intimately connected to assertion. What's wrong with saying things that aren't true? In what way are such assertions defective? If a speaker tells a falsehood deliberately, she is lying, and that might be morally wrong. But not all falsehoods are told deliberately, and it's implausible that telling a falsehood accidentally is always morally wrong. Perhaps some lies are aesthetically ill-advised, but so are some acts of truth-telling. Given (3), however, we have at least a promising outline for an answer to the hard question: we can say that the mistake made by an asserter who violates a norm intimately connected to the speech act is akin to the mistake made by a player who violates either a constitutive or a purposive norm governing the game he is playing.

Unfortunately, promising as these ideas may be, they can't rescue the package consisting of (1)-(3). I shall establish this by defending two major claims. First, even if assertion has constitutive norms, these will be much weaker than any of the alethic or epistemic norms mentioned in (1). Second, depending on how we understand purposes, either the norms generated by such purposes are not intimately connected to baseball (or assertion) after all, or else assertion doesn't have a purpose in the relevant sense. If this is right, then, I suggest, we have sufficient reason to reject (3), and should at least pause before accepting (2).

Two final preliminaries, before proceeding to my main argument: first, it will be useful to have precise formulations of the norms I shall be discussing in this paper. In keeping with the literature, I shall use the following formulations.

(Truth norm)	Assert p only if p .
(Knowledge norm)	Assert p only if you know p .

There are, of course, other alethic and epistemic norms besides the ones just mentioned.⁵ But it will keep my discussion manageable to focus on this pair, and much of what I say will generalize to other such norms (e.g., *Assert p only if you justifiably believe p* , or *Assert p only if you are certain that p*).

Second, in order to keep the discussion manageable, I will mainly concentrate on Williamson's influential defense of the package of views mentioned at the start.⁶ But again, what I say here

⁵ There are some very weak norms that mention truth or knowledge (e.g., *Assert p only if someone knows p* , or *Assert p only if p has a truth-value*). These are not what I have in mind in talking about alethic or epistemic norms, though I shall occasionally mention such very weak norms along the way.

⁶ For discussions of Williamson's account, see DeRose (2002), Hawthorne (2004), Weiner (2005), Douven (2006), Lackey (2007), and Stanley (2007) to mention but a few.

should generalize to other versions of the package, and I'll occasionally mention ways in which the arguments can be used against other defenders of the package.

§ 2 *Constitutive Norms*

Constitutive norms, as we have already seen, are norms that specify what it is to play a game at all, rather than no game, or some other game.⁷ They are to be contrasted with regulative norms, which assess different ways of playing a game: they specify what it is to play the game well, but presuppose that there is something that counts as playing the game in the first place. By way of illustration, consider the following norms of baseball.

- (Three strikes norm) A batter is out when he has earned three strikes.
(Hit batter norm) Pitchers should not hit batters on purpose.

The first of this pair is a constitutive norm of baseball, for it (partly) defines a move in the game, i.e., what it is for a batter to be out. The second is a regulative norm of the game, for it (partly) specifies what counts as playing the game well, by telling us what qualifies as unsportsmanlike conduct within the game.

There is a clear sense in which the constitutive norms of a game are intimately connected to the game, more so than its regulative norms.⁸ This suggests one way of spelling out what it means to say that some norms are more intimately connected to assertion than others: just as moves within games are defined by constitutive norms, perhaps assertion – understood as a move within the larger game of language use – is similarly defined by constitutive norms.⁹ To see if this is plausible, we need to know more about constitutive norms.

The notion of constitutive norms is intended to capture the idea that, for some practices, there is nothing that counts as engaging in the practice “outside the stage-setting” provided by the constitutive norms of the practice (Rawls 1955: 25). Rawls characterized the role that constitutive norms play in making it possible to play a game thus.

To engage in a practice, to perform those actions specified by a practice, *means to follow the appropriate norms*. If one wants to do an action which a certain

⁷ What I am labeling *constitutive norms* are often called *constitutive rules*. I use the term *norm* as a general term that includes rules, leaving open whether there are norms that are not rules.

⁸ It's plausible that the truth and knowledge norms are regulative norms of assertion. But so are sincerity, tactfulness, politeness, relevance, sensitivity (in Nozick's sense), justification, certainty, and any number of other norms, since these are all good-making features of assertions. So, the truth and knowledge norms are no more intimately connected to assertion just in virtue of being regulative norms of the speech act than any of the other norms just mentioned.

⁹ The idea that language use is governed by constitutive norms (or rules) is a familiar one. Williamson, for one, begins his discussion of assertion by stating that his aim is “to identify the constitutive rule(s) of assertion, conceived by analogy with the rules of a game” (2000: 238). In addition, both Searle (1969) and David Lewis (1979) also talk about constitutive rules of language use, understood (in both cases) by analogy to constitutive rules of competitive games.

practice specifies then there is no way to do it except to follow the [norms] which define it (1955: 26; emphasis added).

On this view, then, the role of constitutive norms in baseball can be glossed as follows: in order to play baseball, a player must *follow* the constitutive norms of the game.

Unfortunately, this simple gloss can't be right. Consider the following example. A batter, David, earns three strikes in a particular plate appearance. The umpire miscounts, and instead of ruling that David is out, allows him a fourth strike. (To keep the case simple, assume that no one else present notices the mistake, and that the umpire is not usually prone to such mistakes.) Here, we are inclined to say that David continues to play baseball, even though he fails to conform to a constitutive norm of the game. Since following a norm requires at least conforming to it, the simple gloss must be rejected.

In the example just considered, David's failure to conform to the three strikes norm is both inadvertent and generally unnoticed. Contrast that case with the following modified version. Another batter, Derek, has earned three strikes in a plate appearance. As is usual, the umpire has ruled that Derek is out. However, Derek refuses to leave the batter's box. He admits that he has already earned three strikes, but insists that the game would be more enjoyable for the spectators if he were allowed an extra strike. In this second case, if Derek is sufficiently recalcitrant about being allowed the fourth strike, we are inclined to say that he has stopped playing baseball altogether.¹⁰ The first case (i.e., David's) suggested that inadvertent and unnoticed failures to conform to a constitutive norm are consistent with continuing to play the game; the second case (i.e., Derek's) suggests that intentional and sufficiently marked failures are not.

Let's say that a failure to conform to a norm is *flagrant* if it is intentional and sufficiently marked. In light of the discussion above, the simple gloss considered earlier can be replaced with a weaker principle: in order to play baseball, a player must not *flagrantly* fail to conform to the constitutive norms of the game. This principle returns the correct results in the two cases above. A complete account of constitutive norms would need to say more about what else counts as flagrant failures. For the purposes of this paper, however, this partial characterization will suffice.

Returning now to assertion, suppose the speech act, like baseball, has constitutive norms. Then we can ask: is either the knowledge norm or the truth norm a constitutive norm of assertion? In what follows, I'll offer three reasons to think that the three strikes norm bears a crucially different relationship to baseball than either the knowledge or the truth norm bears to assertion.

First, compare the formulation of the knowledge and truth norms to that of the three strikes norm.

(Knowledge norm)	Assert <i>p</i> only if you know <i>p</i> .
(Truth norm)	Assert <i>p</i> only if <i>p</i> .

¹⁰ The point here is *not* that cheating (or attempted cheating) is always incompatible with continuing to play the game. What is crucial about Derek's case is that his attempted cheating is noticed by the umpires (and everyone else). If a player who is caught attempting to cheat doesn't desist in that attempt, he may cease to play the game.

(Three strikes norm) A batter is out when he has earned three strikes.

If the three strikes norm is taken as a paradigm, then neither the knowledge norm nor the truth norm has the right *form* to be a constitutive norm. The three strikes norm tells us what outs are, what it takes to perform an out at all, not merely at what we should be aiming when attempting to perform one. As such, it has the character of a definition.¹¹ By contrast, neither the knowledge nor the truth norm tells us what it is to assert something. Rather, they each assume that there is something that counts as asserting, and tell us at what an asserter ought to be aiming when performing the speech act.¹²

Second, though there are no doubt some breaches of the three strikes norm in games of baseball, it seems likely that these are extremely rare. For example, by and large, batters are indeed ruled out when they earn three strikes. By contrast, if either the knowledge or the truth norm were a constitutive norm of assertion, then (as Williamson notes) breaches of the constitutive norm of the speech act would be very common indeed (2000: 240).

Third, and most importantly, utterances that constitute flagrant failures to conform to either the knowledge or the truth norms can nevertheless count as assertions. Therefore, both norms violate the principle discussed above, according to which the constitutive norms of a practice are such that it is impossible to engage in the practice while flagrantly failing to conform to them. Consider the case of Pete Rose.¹³ In 2004, Rose admitted that he intentionally told a falsehood when he said several years earlier, “I’m not a chronic gambler” (Crasnick 2004). Even before Rose’s 2004 admission, most baseball fans thought that his earlier utterance expressed a falsehood. So, this was an intentional and highly marked failure to conform to both the truth and knowledge norms. Such failures count as flagrant failures. Nevertheless, surely Rose managed to assert that he is not a chronic gambler. (It is hard to explain why baseball fans were so angry with him for his earlier denials, unless they counted as assertions.) Then, given the principle above, it follows that neither the knowledge norm nor the truth norm is a constitutive norm of assertion.¹⁴

¹¹ Both Searle (1969) and Lewis (1979) emphasize this aspect of constitutive rules.

¹² Can we rewrite either the knowledge norm or the truth norm so that they more closely resemble the three strikes norm with respect to form? Here’s an attempt to rewrite the knowledge norm.

A speaker asserts *p* when he utters a sentence appropriately related to *p* and is criticizable for not knowing *p*.

The phrase *a sentence appropriately related to p* cannot be replaced with *a sentence whose semantic content is p* because it’s possible to assert something without using a sentence that has that as its semantic content. To evaluate this rewritten rule, we need to know what kind of criticizability is at issue here (not to mention what counts as an appropriate relation). But to explain this notion of criticizability is, in effect, to answer what I called *the hard question* in § 1, and answering that question was part of what the analogy with games was supposed to do in the first place. If the analogy cannot do even this, then it’s not clear what philosophical work it’s doing in the account of assertion.

¹³ Pete Rose is a former professional baseball player who received a lifetime ban from the game for gambling on professional baseball. For years, Rose denied the gambling allegations, until the publication of his autobiography in 2004, in which he admitted that he was guilty as charged.

¹⁴ Although Searle agrees with Williamson that assertion is governed by constitutive norms, he (Searle) takes the following to be the constitutive rule of assertion.

An assertion that *p* counts as an undertaking to the effect that *p* represents an actual state of affairs (Searle 1969: 66).

In fact, a much stronger point can be made here. Given the principle above, only norms that are far weaker than the knowledge and truth norms can be constitutive of assertion. After all, it is very easy to make assertions. A speaker can make an assertion even when it is clear to her audience that she does not believe (let alone justifiably) what she is saying. Very young children can make assertions. Therefore, any constitutive norms governing the speech act must be far less demanding than anything we have considered here.

Given these differences, it is tempting to conclude that those who have supposed that either the knowledge or the truth norm is a constitutive norm of assertion must have in mind a different conception of constitutive norms than the one presented in this section. That is certainly possible. But here is some (admittedly defeasible) reason to think that Williamson, for one, has in mind something like the current conception of constitutive norms.

One can think of the knowledge [norm] as giving the condition on which a speaker has the *authority* to make an assertion. Thus asserting *p* without knowing *p* is doing something without having the authority to do it, like giving someone a command without having the authority to do so (Williamson 2000: 257, original emphasis).

Here, Williamson compares a speaker who asserts *p* without knowing *p* to someone who issues a command while lacking the requisite authority. Now, it may sometimes be possible for a speaker to issue a command while lacking the requisite authority, such as when it is mistakenly supposed that she (the speaker) does have the needed authority. But a speaker who *clearly* (i.e., flagrantly) lacks the authority to command can't perform the speech act at all. Consider a student who says to her teacher, "There will be no final exam for this class," or a child who says to his parent, "I will no longer have a nightly curfew." Here, we are not inclined to say that the teacher has been commanded to have no final exam, or that the parent has been commanded to lift the nightly curfew. Thus, it seems that in order to command, a speaker must not flagrantly fail to have the authority to perform the speech act. But then the norm that requires that a speaker who commands have the authority to do so satisfies the principle governing constitutive norms mentioned earlier, according to which the characteristic feature of such norms is that it is impossible to engage in a practice while flagrantly failing to conform to its constitutive norms. So, this norm of commanding is a constitutive norm of the speech act under the very conception of constitutive norms being discussed in this section. And given the parallel between commanding and asserting drawn in the passage quoted above, there is then reason to think that Williamson takes the knowledge norm to be a constitutive norm of assertion under the same conception. If this is right, then there is also reason to think that Williamson's conception of constitutive norms is, after all, closely related to the one under consideration in this section.

Accordingly, we might ask whether the Pete Rose case poses a challenge to Searle's view as well. I think it does, though we need a better grasp of the locution *counts as an undertaking* to be sure. Recall that Rose knew that his original utterance ("I'm not a chronic gambler") was false. Much of his audience thought that the utterance was false. Suppose we allow further that the audience *knew* that the utterance was false. Under these circumstances, it seems to me that it isn't possible for Rose to do anything that counts as an undertaking that the content of his assertion represents an actual state of affairs. So, here again, we have an intentional and marked failure to conform to a candidate constitutive rule of assertion.

In sum: in this section, I have sketched a conception of constitutive norms that takes the three strikes norm as a paradigm. This conception allowed us to characterize a sense in which a norm might be intimately connected to a speech act. And I offered some reason to think that Williamson has something like this conception in mind when he says that the knowledge norm is intimately connected to assertion. But I argued that, under this conception, the constitutive norms of assertion must be much weaker than both the knowledge and truth norms, so this avenue is closed off.

§ 3 *Purposive Norms*

Let's turn now to a different way of unpacking the analogy between language use and games. Here I begin with the observation that playing games like baseball is purposive behavior. Any baseball player has a variety of purposes while playing the game. These may include: whiling away some time, hitting a home run, entertaining spectators, and winning the game, to mention just a few. However, one of these purposes – namely, winning, i.e., scoring more runs than the opposing team – seems more intimately connected to the game than the others, in the sense (it might be held) that it is a purpose that baseball players have *qua* players of the game. That is to say, anyone who plays baseball *must* have winning among their purposes, but need not have any of the other purposes mentioned above. Someone who does not care at all about winning – who, for example, doesn't even attempt to hit pitches, or field ground balls, or anything else – is just not playing baseball; the same need not be said about someone who doesn't care about hitting home runs.¹⁵

If winning is indeed intimately connected to baseball, then moves within baseball – e.g., hitting, base-running, pitching, etc. – can be regarded as having purposes that are *derivatively* intimately connected to them. In each case, the purpose of the move will be to contribute as much as possible, in the manner appropriate to the nature of the move, to achieving the overall purpose of the game, namely, winning.

This suggests a second way of spelling out the analogy between language use and games: given that language use, like baseball, is a purposive activity, perhaps it also has a purpose that is intimately connected to it in just the same way that winning is connected to baseball. Further, just as moves within baseball can be regarded as having purposes that are derivatively intimately connected to them, perhaps assertion, regarded as a move within the larger game of language use, can also be regarded as having a purpose that is derivatively intimately connected to it. Pursuing the analogy, this would be a purpose that any asserter must have *qua* asserter, i.e., in order to count as asserting at all. Like any other purpose, this purpose generates norms.¹⁶ I shall say that norms such as these, i.e., norms that are derived from a purpose that every asserter much

¹⁵ What I have just laid out is one way we might try to spell out what it means to say that winning is intimately connected to baseball. As we shall see (in § 3.1), there are reasons to think that this attempt isn't ultimately successful.

¹⁶ A norm *N* is generated by a purpose *P* iff *N* requires that agents achieve *P*. Since norms are standards for appropriate praise or criticism (see note 1), to say that *N* requires that agents achieve *P* is just to say that agents who are subject to *N* are praiseworthy for achieving *P*, and criticizable for failing to achieve it.

have in order to count as asserting at all, are *purposive norms* of the speech act. Precisely because purposive norms are derived from a purpose that asserters must have *qua* asserter, every asserter must regard her asserting behavior as regulated by these norms.

As I've already noted, for a baseball player, having winning as one of his purposes when playing a game is compatible with having any number of other purposes as well. Moreover, when a player has some purpose that is in conflict with his aim of winning, he may resolve that conflict by subordinating the latter purpose to the former. The same goes, *mutatis mutandis*, for asserters.

One attraction of this approach is that it offers a straightforward answer to what I earlier labeled *the hard question*: what kind of mistake does an asserter make when she violates a norm that is intimately connected to assertion? If that norm is a *purposive* norm of assertion, then the norm must be generated by an aim that every asserter has. Then, if an asserter says something that violates such a norm, she fails in an aim of hers, and thus is subject to practical criticism for failing to achieve one of her *own* ends.

The current approach thus looks promising. But I'll argue that the promise ultimately can't be fulfilled. The difficulties arise when we try to say in more detail exactly what the connection between playing baseball and winning is supposed to be. In the rest of this section, I shall consider two hypotheses about the nature of this connection: first, that the connection consists in winning being an *intrinsic* purpose of baseball, that is, (roughly) a purpose determined just by the kind of game it is; and second, that the connection consists in winning being an *extrinsic* purpose of baseball, that is, (again roughly) a purpose determined partly by facts extrinsic to the game itself. Regarding the first hypothesis, I shall argue that even if language use generally, or assertion in particular, has an intrinsic purpose in the same sense as baseball (and, as we shall see, there are reasons to doubt this), the norms generated by such a purpose wouldn't be either the truth or knowledge norms. Regarding the second hypothesis, I shall argue that although language use generally, and assertion in particular, do have many extrinsic purposes, norms generated by these purposes are not intimately connected with assertion. Thus, though language use does resemble baseball in being purposive behavior, that resemblance doesn't help rescue the package of views with which I started this paper.

§ 3.1 *Intrinsic Purposes*

It is sometimes suggested that winning is connected to baseball not in virtue of the aims of particular players, but simply in virtue of the kind of game it is. To give a label to this idea (which I will attempt to flesh out below), let's say that winning, i.e., scoring more runs than the opposing team, is the *intrinsic purpose* (or *point*) of baseball.¹⁷ By contrast, purposes such as entertaining spectators, or merely whiling away some time, are not intrinsic purposes of the game.

¹⁷ The phrase *intrinsic purpose* is due to Michael Glanzberg (2003). As Glanzberg notes, Dummett (1993) uses the expression *point* for a similar notion. I shall use the expressions interchangeably.

The notion of the point of a game is perhaps not unfamiliar from ordinary (non-philosophical) contexts. For example, to teach someone chess, one might, after explaining how the pieces move, say something like, “The point of the game is to checkmate your opponent’s king.” Similarly, to explain the difference between (regular) chess and variants such as suicide chess, one might say, “The point of chess is to checkmate, while the point of suicide chess is to lose all one’s pieces.” Chess and suicide chess differ in other respects as well: for example, in suicide chess but not in (regular) chess, a player must capture any of her opponent’s pieces that are captureable. Nevertheless, to explain the difference between the games, to explain what sorts of games they are, it isn’t enough merely to describe how their *formal descriptions* differ, i.e., how the games begin, what moves are permissible, and how they end. Besides these differences, it is also necessary to note that the games have different *points*.

I noted earlier that baseball players can play the game with many different purposes. Most of these are optional, in that a player can, but need not, have any given one of them in order to count as playing baseball. But intrinsic purposes might be supposed to differ from other purposes in that, no matter what a player’s other purposes are, he must aim at the intrinsic purpose in order to play the game at all.

This claim might seem subject to obvious counterexamples. Consider, for example, a player who intends to lose a particular baseball game. (Here I have in mind someone who wants to play a complete, but losing, game, not someone who is willing to forfeit the game altogether.) The simplest thing to say about such a player is that he doesn’t have winning among his purposes. But that may be too quick. If the player doesn’t even attempt to hit pitches, or catch fly balls, or run the bases, he runs the risk of ceasing to play the game. In order to achieve his main purpose, he needs to at least try to get some hits, and to catch some fly balls, even if he doesn’t care whether he fails. But to try to do these things is just to try to score runs, and to prevent other teams from scoring runs. And we might say that that means that he’s trying to win, since he’s trying, at least from time to time, to do the things that constitute winning. Thus, we might also say, in order to lose a game of baseball, or for that matter, to achieve any other purpose a player might have, he must try to win the game (though, perhaps, he shouldn’t always try too hard).

I’m not sure that this attempt to unpack the notion of intrinsic purpose ultimately works. Even if it’s true that a player who wants to lose must try to do things that generally lead to winning, if he intends to stop doing those things before he actually wins, then arguably, he isn’t trying to win. If I jog fifty meters, then I’m taking the first few steps towards running a marathon. But if I stop after fifty meters, as I intended to do all along, then I was never trying to run a marathon. The same seems true of the player who intends to lose. That’s bad news for the idea that the intrinsic purpose of a game can be characterized as a purpose that all players must have in order to play the game. And that, I think, ultimately gives us reason to worry about the coherence of the notion of intrinsic purpose, since it is unclear how else to cash out the notion. But I’ll set such worries aside here, because my main aim is to argue that *even if* the notion of an intrinsic purpose is coherent, it doesn’t salvage the package of views with which we started this section.

The point of a game also yields a notion of what counts as good play within the game that is independent from the purposes of the particular players.¹⁸ We can, for example, evaluate a pitch as a good pitch, regardless of the purposes of the players involved. Even if the pitcher in question really wants to lose the game, it can make sense to say that a pitch that fools an opposing batter is a good pitch (and to praise the pitcher for throwing such a pitch). Similarly, we can say to a parent who wants to teach his child to hit but is throwing pitches much too hard for the child, “Those pitches are too good for a child.” When we make such evaluations, we don’t mean that the pitches are good from the point of view of the pitchers’ own purposes. Rather, we make these evaluations based on the contributions the pitches make (or would make) to achieving the point of baseball, even as they may frustrate the pitchers’ primary aims.

Given this notion of the point of a game, we can also derive the points of *moves* within the game. For example, we can say that the point of stealing a base is to increase the probability of scoring a run, where scoring a run increases the probability of scoring more runs than the opposing team, and therefore, of winning the game. This suggests the following schema. If *m* is a move in a game *g*, and the point of playing *g* is *p*, then the point of *m* is generated by what *m* does that, down the line, helps to achieve *p*.

Returning now to the analogy between games and language use, suppose that the latter does have a point, in the same sense that winning is the point of baseball. Then assertion, considered as a move within the game of language use, might also have a point, to be derived from the point of language use in accordance with the schema mentioned in the previous paragraph. And finally, the point of assertion would generate norms which, being purposive norms, could reasonably be said to be intimately connected to the speech act.

Language users, like baseball players, have many purposes when engaging in the activity. Even asserters, a sub-class of language users, have any number of purposes when producing assertions. Thus, an asserter may intend to convey some information to his hearer. Or else, he may intend to guide rational action. Or he may merely intend to express his own beliefs on some matter. And so on. Language users more generally may have even further purposes. A lost traveler may ask for directions. An umpire may call a strike. An enraged manager may express his feelings about the umpire’s eyesight. And so on. But crucially, in language use, unlike baseball, it does not seem that *one* of these purposes is somehow privileged, and thus more intimately connected with the activity, than all the other purposes.

The lack of a privileged purpose in connection with language use might raise the worry that the analogy with games is already strained. But that would be too quick. After all, some games – consider, for example, boxing – have multiple winning conditions, and it might be that language use is like them. Perhaps, then, we should conclude that language use, unlike baseball, has a *disjunctive* point.¹⁹ (Indeed, attention to the variety of purposes that language users have

¹⁸ More carefully, if every baseball player must aim to win, then the claim should be this: the point of a game yields a notion of good play that is independent of the *non-mandatory* purposes of the particular players. For ease of exposition, I shall suppress this qualification from now on.

¹⁹ According to several historically influential taxonomies, speech acts can be categorized according to which of a handful of purposes they aim at. J.L. Austin (1975), Searle (1979), and Bach and Robert Harnish (1979) all agree on this, though they disagree about what the purposes in question are. Searle, for example, thinks that the purposes are

suggests that that point would have to be *deeply* disjunctive.) Winning in baseball would then correspond to achieving *one* of the disjuncts that constitute that disjunction.

But there are still deeper problems here. As we shall see next, even if there is such a disjunctive point for language use, there are still crucial disanalogies between that point and the points of games like baseball. (It doesn't follow, of course, that language use has no intrinsic purpose, in some sense of *intrinsic purpose*, but only that that claim derives no support from the analogy with games.) Further, even if language use did have such a disjunctive point, there would still be trouble for the claim that either the knowledge or truth norm is intimately connected to assertion.

First, if the point of language use is disjunctive, in fact, deeply disjunctive, then it will be easy, almost trivial, to achieve that point when performing a speech act. If that point includes disjuncts that are meant to account for the variety of purposes mentioned earlier, then, when a speaker performs a speech act, she will ordinarily satisfy at least one of the disjuncts (even if not the disjunct she intended to satisfy).²⁰ It will be difficult, then, to perform a speech act while not satisfying the point of language use. By contrast, it is not at all difficult to play baseball while failing to achieve its point; one of our two teams does so in every game played. Indeed, it's a virtue of game design that their points are non-trivial to achieve.

Second, and more importantly, the intrinsic purpose of language use, unlike that of baseball, doesn't yield a notion of good play that is independent of the purposes of particular language users. As I've already noted, we can, and do, judge moves in baseball as good or bad independent of the particular purposes of the players. But we rarely, if ever, judge speech acts independently of the purposes of particular language users. We don't say that promises, suggestions, warnings, and so on are good even when they frustrate the purposes of the speakers and hearers involved. For some speech acts, including assertions, it's not even clear that we have the locutions for praising them independently of the purposes of language users. ("Good assertion!" and "Well asserted!" both sound awkward, as does "You know that!". "Well said!" is generally used to praise style. "That's true!" is not a clear indicator of praise.) Finally, it's hard to make sense of admonitions like, "You shouldn't give warnings that good in these circumstances," or "You shouldn't make assertions that good right now." (Compare "Those pitches are too good for a child", discussed earlier.) We do sometimes say things like, "You shouldn't ask questions that good in an undergraduate talk," but by this, we don't mean that the questions are good from the point of view of the intrinsic purpose of questioning (or language use). Rather, we generally mean that the questions are too hard for the person being questioned, say, an undergraduate speaker.

these: committing the speaker to something's being the case; getting the hearer to do something; committing the speaker to a future course of action; expressing certain psychological states (Searle 1979: 12-5). On all these views, there is no one privileged purpose associated with language use, but a few different purposes. Thanks especially to an anonymous referee for Oxford University Press for help on this point.

²⁰ Things are more complicated if each speech act has exactly one of the disjuncts as its intrinsic purpose. In that case, it may not be so trivial to achieve that purpose when performing the act. But then, there's a distinct way in which the intrinsic purposes of language use and games differ: if each speech act has its own intrinsic purpose, and the intrinsic purpose of language use is just the disjunction of those purposes, then it seems that the intrinsic purpose of the latter is derived from the intrinsic purposes of its constitutive moves, rather than the other way around (as in the case of games).

Third, even if language use does have an intrinsic purpose in a sense analogous to baseball, it still doesn't follow that either the truth or knowledge norms is intimately connected to assertion. If anything, close attention to the analogy makes trouble for these claims. Recall that the point of any move within a game is given by what that move does that, down the line, helps to achieve the point of the game. Notably, that means that a move can achieve its point, and thus be a good move, even if it doesn't, by itself, bring about the win. Consider, for example, pitching in baseball. A pitch can achieve its point even if it doesn't itself bring about the win, or even an out. In fact, a pitch can be a good pitch even if it doesn't constitute a strike, if it just raises the probability that the next pitch (or one later on) will be a strike. That is to say, whether a particular pitch achieves the point of the move depends on several things, including the circumstances in which the pitch is thrown, and what happens in the game subsequently. The same pitch can be a good pitch in one game situation, and a terrible one in another.²¹

Pursuing the analogy, we should say that whether a speech act achieves its point depends not merely on factors intrinsic to the act itself, but also on other factors, including the circumstances in which it's produced, and the speech acts that follow. This will be true regardless of what the point of language use turns out to be, whether it includes transmitting knowledge, guiding rational action, truly describing the world, or anything else. Applied to assertion, this means that whether an assertion achieves the (derivative) point of the speech act – i.e., whether it does something that, down the line, helps achieve the point of language use, whatever that may be – depends on various factors extrinsic to the assertion itself. And it seems implausible to suppose that only assertions that are true (or known) can, down the line, achieve the point of language use.

Here's an example to make this vivid. Suppose for a moment that part of the point of language use is to transmit knowledge. (Williamson briefly suggests that this is the point of *assertion* (2000: 267). It certainly seems a friendly suggestion for any defender of the knowledge norm.) Then, consider the extended conversation about norms of assertion of which this paper is a part. Many of the papers that have been part of this conversation have been excellent, not in any way defective. And that's true even though several of those papers disagree with each other, and so, cannot contain only true claims. No matter. If any of those papers were the last word on the question, not containing only true (or known) claims might well be a flaw. But that's not what academic papers, or in general, assertions in academic conversations, aim to do. Rather, they try to move the debate along, to raise new suggestions, to perhaps point out inconsistencies or mistakes that have crept into the conversation. It doesn't matter for these purposes whether all the assertions they contain are true. What matters is that they help us, collectively, drive towards knowledge. In other words, academic conversations can, and do, achieve that aim, even when some of the assertions that make up the conversation are false. Even more strongly, requiring that all assertions that make up such conversations be true (or known) might well hinder the aims of accumulating and transmitting knowledge.

²¹ The same points can be made using other games, such as chess. Clearly, a move in chess can be a good move even if it doesn't checkmate. Rather, a move is good if, down the line, it contributes to winning the game. Further, it's generally good play in chess to save one's own pieces, and capture one's opponents'. But in some circumstances, sacrificing a pawn is a good move, especially if it garners some advantage in later play. Thus, what constitutes a good move is importantly context-dependent.

In the previous paragraph, I assumed, for the sake of argument, that transmitting knowledge was one of the disjuncts that constituted the point of language use. But similar examples can be constructed for other candidate disjuncts. Overall, the key point is that, just as a move in baseball can achieve its point as long as it contributes, down the line, to winning, even if that move doesn't, by itself, bring about the win, so an assertion can achieve its point as long as it contributes, down the line, to achieving the point of language use, even if it doesn't, by itself, bring about that point. The only way to derive the knowledge or truth norms in this context is by insisting that good assertions must be the equivalent of winning moves, i.e., moves that, by themselves, bring about wins. But that idea is not at all supported by the analogy with games. What that analogy in fact suggests is something quite different, namely, that good assertions are those that lead, perhaps a long way down the track, to knowledge or truth (or whatever else the point of language use might be).

To close this discussion, let me emphasize one aspect of the analogy with games that needs to be treated with care. As has been emphasized by many philosophers, language use in general, and conversations more specifically, are cooperative endeavors, not competitive. Whatever the point of language use, it can be achieved by all the participants simultaneously, working together. Conversations are not zero-sum. This does not mean that the analogy with games is mistaken. But it does mean that it might be better to analogize language users to, say, teammates in baseball, rather than, as is often done, to opponents in chess or some other competitive game.

Before ending this section, I briefly want to consider an objection to my argument. Thus far, I have been focusing on a picture according to which assertion, regarded as a move within the larger game of language use, can be seen as derivatively having a point. But, it might be objected, this complicates things unnecessarily. Instead of likening language use to games, why not liken *assertion* to games, and regard assertion as having a point non-derivatively, as games do? Perhaps, the objection continues, we could say that the point of assertion is to transmit knowledge, or to guide rational action.

In fact, this move doesn't change things much. Several of the points made above still apply. In particular, it still seems that the point of assertion, unlike the point of games, doesn't yield a notion of good play. That doesn't mean that assertion doesn't have a point, in some sense of *point*, but it does suggest a disanalogy with baseball. Further, even if one of the above candidates *is* the point of assertion, it still doesn't follow that every good assertion must be known or true, just as it doesn't follow from the fact that winning is the point of baseball that every good baseball play must be a winning play.

It's worth pointing out that some weaker claims about the point of assertion seem *prima facie* at least as plausible as the claims mentioned above. For example, we could say that the point of the speech act is to say something truth-conditional, or, if we are working within the Stalnakerian framework, that the point of assertion is to make a well-defined change to the context set (Stalnaker 1978). It's plausible, at least, that a speaker who doesn't aim to say something truth-conditional doesn't count as asserting at all. And if a speaker intends to assert something while lacking the intention to say something truth-conditional, he may well be subject to at least practical criticism. But if one of these weaker claims about the point of assertion is correct, then

again, it is hard to see how that point can generate either the truth or knowledge norms, though it may well generate a much weaker norm, such as *Aim to say only what is truth-evaluable*.

§ 3.2 *Extrinsic Purposes*

In § 3.1, I said that intrinsic purposes are connected to games just in virtue of the kinds of games they are. But purposes can also be connected to games (partly) in virtue of facts extrinsic to the games themselves, such as facts about the players (and spectators) of the game. Let's say that these are *extrinsic purposes* of the game. In this section, I shall consider the hypothesis that winning is an *extrinsic* purpose of baseball.

There are many extrinsic purposes associated with any game. There are the purposes that the game serves for us, its players and spectators. There are the purposes that help explain why we care about or enjoy the game, why we have the game in the repertoire of games that we play. There are also the purposes that (all, most, or just some) players have when playing the game. All of these are extrinsic purposes of the game, for they are all purposes that are associated with the game in virtue of facts about its players (and spectators).

Properly speaking, extrinsic purposes should be indexed to groups of individuals. The purposes that help explain why *we* care about baseball might be quite different from those that figure in the explanations for why others care about baseball. Similarly, the purposes that baseball serves for us might also be quite different from those that it serves for others. As such, the extrinsic purposes of baseball with respect to us may be very different from the extrinsic purposes of the game with respect to some other groups. Though I will suppress this indexing for the most part, it will be made explicit where it is relevant below.

Given this (quite expansive and heterogeneous) notion of an extrinsic purpose, there will be many extrinsic purposes of baseball. There are, after all, many purposes that are shared by baseball players. For example, it is surely the case that many baseball players care about hitting home runs. That would make home run hitting an extrinsic purpose of baseball. For us, baseball serves the purpose of entertaining spectators. That makes entertaining spectators another extrinsic purpose of baseball. And so on for many other purposes.

In light of the expansiveness of the notion, it's not surprising that winning would be another extrinsic purpose of baseball. After all, it's plausible that we care about baseball because we care about the effort of trying to score more runs than the opposing team. Surely most baseball players, if not all, have winning as their primary aim when playing the game. Either of these putative facts is sufficient to make winning an extrinsic purpose of baseball. Further, winning being an extrinsic purpose of baseball is compatible with its also being its intrinsic purpose. To keep things simple, however, I will assume for the remainder of this section that winning is *only* an extrinsic purpose of the game.

As in the case of points, besides the extrinsic purposes of games, we can also talk about the extrinsic purposes of moves within games. In fact, we can adopt the very same schema that was discussed in the previous section: if *m* is a move in a game *g*, and *p* an extrinsic purpose of *g*,

then an extrinsic purpose of m is generated by what m does that, down the line, helps to achieve p . Thus, if entertaining spectators is an extrinsic purpose of baseball, then hitting a home run in a way that contributes to the entertainingness of the game may be among the extrinsic purposes of batting.

Returning to the analogy between language use and games, surely language use has extrinsic purposes in the sense at issue here. These would be purposes that are associated with language use in virtue of facts about language users. In § 3.1, we have already mentioned several purposes that language users can have: transmitting knowledge, guiding rational action, expressing emotions, rendering verdicts, asking questions, and so on. All of these, and many others, will be extrinsic purposes of language use.

Extrinsic purposes, like purposes of any other kind, generate norms. Unfortunately, as I will argue below, the norms thus generated cannot be said to be intimately connected to the activity they govern.

First, since language use, like baseball, has many extrinsic purposes, we should expect that these purposes will generate a variety of norms governing language use. All such norms will be equally closely connected to the activity of using language. In other words, none of these norms will be particularly intimately connected to the activity. Even assertion will have many extrinsic purposes, derived from the various extrinsic purposes of language use. Each of these purposes will generate many norms governing the speech act. Again, all such norms will be equally intimately connected to the speech act.

It may be objected that though there are many extrinsic purposes of language use, some of these purposes are special, and it is the special purposes that generate norms that are intimately connected to language use, and to assertion. Which are the special extrinsic purposes? It's possible to imagine several different ways of carving out special extrinsic purposes. For example, we could focus on those extrinsic purposes, if there are any, that are shared by all language users (or all asserters), or we could focus on those purposes that language use (or assertion) serves for *us*, and so on. But the main point to emphasize here is that the analogy with games gives us no guidance with respect to privileging certain extrinsic purposes over others. If there are some privileged extrinsic purposes associated with language use, or with assertion, the analogy with games does no work in pointing us toward them.

Second, the norms generated by extrinsic purposes are not purposive norms. Recall that purposive norms for any activity are norms that are generated by purposes that everyone must have in order to participate in that activity. But extrinsic purposes do not generally have this property. Suppose again that home run hitting is an extrinsic purpose of baseball because it is a purpose shared by many baseball players. Even if this is true, it needn't be true that *everyone* aims to hit home runs. But those who don't care at all about hitting home runs can nevertheless play the game. Thus, the norms generated by the purpose of hitting home runs aren't purposive norms.

I said earlier that one of the reasons to focus on purposive norms is that they provide a straightforward answer to what I called *the hard question*, namely, the question about what kind

of mistake someone makes when they fail to satisfy a norm intimately connected to an activity. If someone intends to play baseball, but violates a purposive norm of the game, he will not count as playing the game. As such, he is subject to practical criticism for failing to achieve one of his *own* ends. However, if the norms in question aren't purposive, it is much harder to see what sort of mistake someone makes in failing to satisfy them. Consider Alex, a baseball player who likes to play the game but doesn't much care about hitting home runs. If Alex doesn't attempt to hit home runs when playing the game, what kind of mistake does he make? On what grounds can we criticize him? He makes no mistake by his own lights. He may frustrate *our* desires to see home runs, but we can't generally criticize others whenever they frustrate our desires. For analogous reasons, when someone violates a norm governing assertion, if that norm is merely generated by an extrinsic purpose, it is not all clear that any answer to the hard question will be forthcoming.

Here it may be objected that though the extrinsic purposes of baseball could have been quite different from what they in fact are, that's not true for assertion (and for language use more generally). In particular, it might be argued, we care about language use in part because we care about transmitting knowledge, but given the kinds of limited beings we are, we couldn't have done otherwise. As such, the objection continues, transmitting knowledge is an extrinsic purpose of language, and it couldn't have been otherwise. If this is true, however, this points to an important respect of *disanalogy* between games and language use, for all games could survive changes in their extrinsic purposes. The analogy between games and language use, once again, offers no help at the point.

In sum: in the last two sections, I have sketched conceptions of intrinsic and extrinsic purposes. I've raised some concerns about the coherence of the notion of intrinsic purpose as applied to games, and about the applicability of the same notion to language use. And I've argued that the norms generated by the intrinsic purpose(s) of language use are much weaker than the knowledge and truth norms. By contrast, there are no analogous worries about the coherence of the notion of extrinsic purpose, either as applied to games or to language use. But I've offered reasons to think that the norms generated by extrinsic purposes aren't intimately connected to the activities they govern.

§ 4 Conclusion

To conclude, let's return to the package of views with which I started this paper. Recall that that package consisted of three views: first, that assertion is governed by either an alethic or an epistemic norm; second, that this norm is more intimately connected to assertion than other norms governing the speech act; and third, that this sense of intimate connection can be understood with the help of the analogy between language use and games. I have been arguing here that the third view in this package should be rejected, that although language use and games may be analogous in some ways, those respects of analogy don't help us get at a sense in which some alethic or epistemic norm is more intimately connected to assertion than other norms.

My conclusion is thus in a way quite modest. In particular, I have not shown that there is no way of making sense of the second view in this package, and I have certainly not cast doubt on the

first view there. My main point here is not that assertion (and other speech acts) are not governed by norms, or even that assertion in particular is not governed by an alethic or epistemic norm, but rather, that we need a better understanding of the relationship between norms and the speech acts they govern. Insofar as we have been relying (explicitly or tacitly) on the analogy with games to make sense of this relationship, as I suspect we have, I hope that this paper offers reason to think that we need to look elsewhere.²²

References

- Austin, J.L. 1975. *How to Do Things with Words*, Second Edition, J.O. Urmson and M. Sbisá (eds.). Cambridge, MA: Harvard University Press.
- Bach, K. Forthcoming. Knowledge in and out of Context. In Joseph K. Campbell and Michael O'Rourke (eds.), *Knowledge and Skepticism*. Cambridge, MA: MIT Press.
- Bach, K. and R. Harnish. 1979. *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.
- Crasnick, J. 2004. Covering Rose Never Dull. *ESPN Insider* (<http://proxy.espn.go.com/mlb/columns/story?id=1702424>).
- DeRose, K. 2002. Assertion, Knowledge, and Context. *The Philosophical Review* 111: 167-203.
- Douven, I. 2006. Assertion, Knowledge, and Rational Credibility. *Philosophical Review* 115: 449-85.
- Dummett, M. 1978. Truth. In his *Truth and Other Enigmas*, 1-24. Cambridge, MA: Harvard University Press.
- Dummett, M. 1993. What is a Theory of Meaning? (II). In his *The Seas of Language*, 34-93. Oxford: Clarendon Press.
- Glanzberg, M. 2003. Against Truth-Value Gaps. In J.C. Beall (ed.), *Liars and Heaps: New Essays on Paradox*, 159-94. Oxford: Oxford University Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hindriks, F. 2007. The Status of the Knowledge Account of Assertion. *Linguistics and Philosophy* 30: 393-406.
- Lackey, J. 2007. Norms of Assertion. *Noûs* 41: 594-626.
- Lewis, D.K. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic* 8: 339-59.
- Price, H. 1998. Three Norms of Assertibility, or How the MOA Became Extinct. *Philosophical Perspectives* 12: 241-54.
- Price, H. 2003. Truth as Convenient Friction. *Journal of Philosophy* 100: 167-90.
- Rawls, J. 1955. Two Concepts of Rules. *Philosophical Review* 64: 3-32.
- Searle, J. 1969. *Speech Acts: An Essay in the Philosophy of Language*. London: Cambridge University Press.
- Searle, J. 1979. *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge: Cambridge University Press.
- Stalnaker, R. 1978. Assertion. *Syntax and Semantics* 9, 315-32. New York: Academic Press.
- Stanley, J. 2007. Knowledge and Certainty. *Philosophical Issues* 18: 33-55.
- Weiner, M. 2005. Must We Know What We Say? *Philosophical Review* 114: 227-51.

²² For helpful comments on various earlier drafts of this paper, I'd like to thank Mary Kate McGowan, Tom McKay, Brian Montgomery, Brian Weatherson, audiences at the University of Rochester, University of Western Ontario, and Queen's University, as well as an anonymous referee for Oxford University Press.

- Williamson, T. 2000. Assertion. In his *Knowledge and Its Limits*, 238-69. Oxford: Oxford University Press.
- Williamson, T. 1996. Knowing and Asserting. *Philosophical Review* 105: 489-523.
- Wright, C. 1992. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.