

Truth and assertion: rules versus aims

NERI MARSILI

1. A puzzle about assertion and truth

In his seminal paper “Knowing and Asserting” (1996/2000), Williamson proposes to analyse under which conditions an assertion is proper *qua* *assertion*. He presents a simple hypothesis:

An attractively simple suggestion is this. There is just one rule. Where C is a property of propositions, the rule says: [...] “One must: assert p only if p has C.”

To identify the norm of assertion, then, one needs to specify property C. Several proposals have emerged in the literature. Broadly, they can be grouped into four categories:

- (KR) Knowledge-Rule: “Assert p only if you know that p”
Williamson (1996/2000), DeRose (2002), Hawthorne (2004)
- (TR) Truth-Rule: “Assert p only if p is true”
Alston (2000), Weiner (2005), Whiting (2012)
- (JR) Justification-Rule: “Assert p only if you rationally believe that p”
(or “only if it is rational for you to believe that p”)
Douven (2006), Lackey (2007), Kvanvig (2009)
- (BR) Belief-Rule: “Assert p only if you believe that p”
Hindriks (2007), Bach (2010)

The most demanding of these proposals is KR. Knowledge is understood to entail belief, justification and truth, so that in order to follow KR, you also have to follow all the other rules (JR, BR, TR). Some authors understand JR as requiring you to follow BR, whereas no author takes BR or TR to entail any other rule¹.

¹ Hindriks (2007) and Bach (2010) believe that KR can be *derived* from BR, on the assumption that knowledge is required for proper belief. Unlike cases involving

In this paper, I shall discuss what is arguably the starkest matter of disagreement in this debate: the question of whether a proper assertion needs to be true. Consequently, I will divide these four accounts into two supergroups: those requiring truth for proper assertion (TR, KR) and those not requiring it (BR, JR). I call the former accounts *FACTIVE*, and the latter *NON-FACTIVE* (since they do, or do not, require speakers to only state facts):

FACTIVE ACCOUNTS: KR, TR

NON-FACTIVE ACCOUNTS: BR, JR

Both families of accounts purport to explain a different set of equally plausible, but apparently incompatible, linguistic data. The main linguistic datum supporting factive accounts is what I call *IMPROPER FALSITY*². People share the *intuition* that (*ceteris paribus*) a false assertion is incorrect and improper. This intuition is often translated into action: people typically criticise false assertions in virtue of their being false. *IMPROPER FALSITY* can thus be spelled out both as an intuition and as a behaviour:

IMPROPER FALSITY

- INTUITION: *false assertions are improper in virtue of their being false*
- BEHAVIOUR: *false assertions are reproachable in virtue of their being false*

Factive accounts can explain *IMPROPER FALSITY* in terms of the violation of a factive norm, whereas non-factive accounts cannot (cf. Williamson 2000:262). To be sure, in a subset of cases non-factive accounts are able to acknowledge the improperness of false assertions, given that *some* false assertions violate non-factive rules (JR, BR) – but this only helps in a subset of cases. More importantly, it seems that falsity constitutes a distinctive kind of wrongness for assertion. Any account failing to acknowledge that false assertions are incorrect and criticisable *in virtue of their being false* misses a fundamental linguistic datum about assertion.

entailment, here the source of normativity differs – agents are subject to KR *qua* *believers*, rather than *qua* *assertors*.

² I am not claiming that this is the *only* datum supporting factive accounts. For instance, Williamson (2000) famously presents three further data points in favour of KR: ‘Moorean assertions’, ‘conversational challenges’ and ‘lottery assertions’. I leave these arguments aside because I am concerned with data that supports factive accounts as a whole, as opposed to specific accounts (KR, TR, etc.). Furthermore, it is debatable whether data supporting specific accounts can also be explained by other accounts (cf. Lackey 2007, Kvanvig 2009, McKinnon 2013), whereas there seems to be consensus that factive accounts have a vantage point to explain *IMPROPER FALSITY*.

Proponents of non-factive accounts deny the accuracy of IMPROPER FALSITY. They point out that *inadvertently* saying something false (*i.e.* making an ‘unlucky assertion’) does not generally amount to bad linguistic behaviour, nor does it elicit criticism. To see this, consider the following example.

Francesca has a collection of rocks, that she inherited from her grandfather. The collection has always been kept in a very secure safe. One evening, Francesca is enjoying an aperitif with her friend Baba, who asks her whether she possesses any meteorites. Since one of the rocks in the collection is a meteorite, Francesca replies:

(1) Yes, I have a tiny meteorite in my safe at home

Unbeknownst to Francesca, however, some thieves have entered her house that very afternoon, and stolen all her rocks from the safe.

In uttering (1), Francesca is saying what she reasonably believes to be true: she is following non-factive rules (JR, BR) but violating factive ones (TR, KR), since (1) is false. Non-factivist philosophers share the intuition that it would be inappropriate to deem Francesca’s linguistic behaviour impermissible, or to reproach her for having asserted (1). More generally, they contend that *some* false assertions (‘unlucky assertions’) are permissible. To sum up, their view purports to explain a different set of linguistic data:

PERMISSIBLE FALSITY:

- *INTUITION: inadvertently false (‘unlucky’) assertions are permissible*
- *BEHAVIOUR: we do not reproach unlucky assertors*

Here is the puzzle that this paper addresses. Both IMPROPER FALSITY and PERMISSIBLE FALSITY appear to be plausible descriptions of our linguistic intuitions and behaviour. Only factive accounts are able to explain IMPROPER FALSITY, and only non-factive accounts are able to explain PERMISSIBLE FALSITY. But these two sets of linguistic data (and the corresponding accounts of assertion) seem incompatible: according to the former, false assertions are always improper; according to the latter, they are permissible under some circumstances.

In this paper, I propose a novel solution to this puzzle³. Contrary to what is commonly taken for granted, IMPROPER FALSITY need not be explained in terms of a violation of the *norm* of assertion: it can also be explained in terms of a failure to meet the *aim* of assertion. I argue that truth should indeed be understood as the aim, rather than the rule, of assertion. The resulting account yields a non-factive explanation of IMPROPER FALSITY that still allows for PERMISSIBLE FALSITY, *i.e.* a solution to the puzzle under consideration.

2. Rules vs Aims

The view that assertions are governed by norms is typically conceived “by analogy with the rules of a game” (Williamson 2000:239, Kölbel 2010). An often-undervalued point is that games not only have rules, but also aims – aims that a player purports to have in playing the game (Schwyzer 1969, Maitra 2011). And there is an intuitive difference between the rules of a game and the aims of a game.

Arguably, the essential ‘aim’ of competitive games is that of winning (Kemp 2007:113): in this sense, checkmate is the aim of chess. The difference between aims and rules is quite intuitive here: even if the aim of chess is to checkmate, it is not a rule of chess that you should only make moves that lead you to checkmate. Now, often competitive games also involve derivative, intermediate aims – things that you standardly aim to do in order to win the game. For instance, the aim of shooting a penalty in football is to score a goal, because only scoring a goal will lead you to win the game. These intermediate aims are also not rules: while scoring is the purported aim of shooting a penalty, it is not a rule of football that you should shoot the penalty only if you score a goal.

These examples show that there is an intuitive difference between the aims and the rules of a game. To offer a broad characterisation of this difference, we could define them as follows:

For every condition X, action-type φ , and action-token $\varphi 1$ (of the action-type φ):

- X is the (only) aim of φ iff $\varphi 1$ is successful *if (and only if) X*
- X is the (only) rule of φ iff $\varphi 1$ is permissible (*if and*) only if X

³ An alternative solution would be to deny that this is a genuine puzzle, and follow the factivist philosophers who believe that their view is consistent with PERMISSIBLE FALSITY. I consider and dismiss this alternative strategy in the closing paragraphs of this paper.

Let us consider two examples from football to illustrate this. Saying that ‘scoring a goal’ (X) is the aim of ‘shooting a penalty’ (φ) means that any given instantiation of ‘penalty-shooting’ (φ) in football is successful if a goal is scored (*i.e.* condition X is met). By contrast, saying that ‘players cannot touch the balls with their hands’ (X) is a rule regulating ‘ball-touching’ (φ) in football means that any given instantiation of ‘ball-touching’ (φ) in football is permissible only if the ball is not touched with the hands (*i.e.* condition X is met). The obvious difference is that the aim identifies a sufficient condition for *successful* action, whereas the rule identifies a necessary condition for *permissible* action.

One important similarity between rules and aims, by contrast, is that they can both ground evaluative judgments in the right circumstances. We can criticise agents not only for violating rules, but also for failing to fulfil aims. For instance, supporters may criticise a player for failing to meet aims such as scoring a penalty, or completing an easy pass. This suggests that both aims and rules can ground criticisms and challenges.

If these observations are correct, there are two ways to account for IMPROPER FALSITY: a factive explanation (truth is the *rule* of assertion), and an explanation that is compatible with non-factive accounts (truth is the *aim* of assertion). The latter explanation has an advantage over the former: it enjoys the same explanatory power (it can explain why false assertions are incorrect and criticisable, and why speakers are generally expected to retract them), while avoiding the problem of regarding unlucky assertions as impermissible, and unlucky assertors as reproachable. But even if the truth-aim account fares better in this respect, it is yet to be proved that truth is the aim of assertion.

3. *Is truth the rule or the aim of assertion?*

*There are a thousand ways of missing the bull's-
eye,*

only one of hitting it.

Montaigne, On Liars (*Essays*, 1.9)

If we are to decide whether truth is the rule or the aim of assertion, it would be useful to have a test to tell apart rules from aims. My proposed definitions (of rules and aims) are not helpful here: they presuppose some knowledge of what makes an assertion ‘successful’ as opposed to ‘permissible’, which is exactly what we are trying to establish. A useful test should rely on differences that we can spot just in virtue of our familiarity with the relevant game or practice. One such difference is the following: assuming a competent and careful agent, aims characteristically allow for unintentional failure, whereas

rules do not typically allow for unintentional failure. In other words, typically aims are such that you can try but fail to meet them, whereas rules are such that you cannot easily try but fail to follow them.

To see this, consider some examples of aims: in shooting a penalty, I can fail to score a goal despite my best effort; in playing chess, I can fail to checkmate even if I play to the best of my capabilities. The same is not true for rules: it is difficult for a *competent and careful* player to try but fail to move the bishop only diagonally, or to try but fail to castle only if king and rooks have not moved yet. If these observations are on the right track, the following is a plausible test for distinguishing rules from aims:

TEST FOR THE RULE-AIM DISTINCTION:

For any given action A within a practice P, for any given condition C applying to A, for any competent and careful agent S, the truth of (T) is prima facie evidence that C is the aim of A in P, whereas the falsity of (T) is prima facie evidence that C is the rule of A in P:

(T) S can typically try but fail to meet C

Is the test valid? Let us consider some apparent counterexamples. One is the '8-second rule' in basketball. According to this rule, a team who gains control of the ball in the backcourt "must cause the ball to go into its frontcourt within 8 seconds" (FIBA rulebook 2014:30). It seems that players who are *carefully* trying to follow this rule can sometimes violate it, against the predictions of the TEST. However, this worry dissipates once we realise that the notion of 'carefulness' that I am invoking in the TEST is slightly more demanding.

By 'careful player' I mean a player that prioritises following the rules of the game over trying to meet its aims. Allegedly 'careful' players who violate the '8-second rule', by contrast, typically fail to follow the rule because they prioritise their aims (keep possession of the ball, execute a good pass, etc.) over what is prescribed by the rule. Crucially, these players always have an available alternative course of action that would minimise or completely avoid the risk of violating the rule: perform a (riskier) pass, shoot, or (as paradoxical as it may seem) let the other team gain control of the ball. In sum, rather than a counterexample to the test, this is an example of a rule that is often violated because it conflicts with the aims of the game.

A different sort of counterexample is represented by the rule against touching the ball with your hands in football. Since it is possible to touch the ball with your hands inadvertently, the rule seems to constitute a counterexample to the TEST. However, the actual football rule only prohibits "a deliberate act of a

player making contact with the ball with the hand or arm” (IFAB rulebook 2016/2017:82), so that you cannot try but fail to follow it.

This rule is particularly telling, because it regulates an action that can be easily performed inadvertently. Given the risk of inadvertent hand-touch, a rule prohibiting each and every hand-touch would be unfair – as it would often cause players to be punished for actions they performed inadvertently. To avoid this, the rule is purposefully designed to punish only deliberate violation. I am willing to speculate that a similar point applies to most rules: whenever there is a relevant risk of inadvertent violations, rules tend to be designed to prohibit only intentional violations, because the opposite would be unfair.

A last point to note is that the TEST only specifies a condition for *prima facie* evidence that C is the aim of A: for my purposes, it can allow for some countervailing considerations. Even if the TEST does not provide a conclusive proof, it tips the balance in favour of the truth-aim account. The truth-*rule* account atypically and unfairly allows for inadvertent violations. The truth-*aim* account, by contrast, typically and fairly counts unlucky assertions as unsuccessful but permissible. Furthermore, and relatedly, only the truth-aim account is compatible with PERMISSIBLE FALSITY.

4. The truth-aim account

My proposal to understand truth as the ‘aim’ of assertion is still quite vague, as it is not obvious in which sense a speech act like assertion can aim at a truth-value. My view is that ‘aiming at truth’ means that an assertion is successful only if it is true, because assertions conventionally have a *word-to-world* direction of fit. I will explain both claims, and outline some connections with existing scholarship.

The first claim follows from my proposed account of aims, according to which “X is the (only) aim of $\varphi = \varphi_1$ is successful if (and only if) X”. On this conception, saying that assertion aims at truth means that an assertion is successful (*qua assertion*) if it is true, more or less in the same way in which saying that the aim of a penalty is to score a goal means that the penalty is successful if you score a goal.

Clearly, the relevant standard for successfulness here is the aim set by the relevant practice, not by the agent. In playing chess, I may aim to bore my opponent, thereby preferring a stall to a victory. In succeeding in stalling, I would be successful with respect to my personal aim, but not with respect with the aim of playing chess. Similarly, in asserting I might aim to say something false, and be successful with respect to that aim, but not with respect to the purported aim of asserting. Let us call the aims that are determined by a

practice the *purported* aims⁴ of that practice, and any divergent aim of the agent a *personal* aim. The truth-aim account takes truth to be the *purported* aim of assertion, and sets no direct constraints on the *personal* aims of assertors.

As for the second component of my explanation, that assertions have a *word-to-world* direction of fit is anything but a controversial claim: it is rather the current orthodoxy in speech act theory (e.g. Searle 1969, Recanati 1987:147-63, Vanderveken 1990, Green 2013), and more generally in philosophy of language and linguistics. It is also standard to read the notion of direction of fit as having a teleological connotation, *i.e.* as setting a condition of success for the speech act, as opposed to a condition for permissible performance (Humberstone 1992). As a matter of fact, often authors treat ‘having truth as an aim’ and ‘having a *word-to-world* direction of fit’ as synonymous: “‘true’ and ‘false’ are the favourite terms of assessment for achieving *success* of illocutions which have the word-to-world direction of fit” (Searle 2007:34).

One may object that, while aim-directed moves in a game (such as ‘scoring a penalty’) can be ordinarily described as successful or unsuccessful, assertions do not fit this category of assessment: we do not ordinarily criticise false assertions *as unsuccessful*. However, this objection is based on a misconstrued analogy. ‘Failed in asserting’ is different from ‘failed in *scoring* a penalty’, in that the former expression inappropriately describes as unsuccessful the action itself (asserting), whereas the latter expression appropriately describes as unsuccessful the attempt to meet the aim (scoring). This worry disappears when the analogy is construed appropriately: we can criticise a speaker by saying that she ‘failed to say the *truth*’; relatedly, we would not criticise a player who failed to score a goal by saying that she ‘failed to *shoot a penalty*’. To be sure, I am not claiming that *only* assertions have word-to-world direction of fit. Several speech acts aim at truth, namely all speech acts belonging to the class of ‘assertives’ (also known as ‘representatives’ in Searle 1976, or ‘constatives’ in Recanati 1981): conjecturing, hypothesising, objecting are but a few examples. The success condition for all these speech acts is that their propositional content be true: conjectures, hypotheses, objections that *p* are also successful if *p* is true. The bottom line is that purportedly aiming at truth is a necessary, but not sufficient, condition for an utterance to be an assertion: the truth-aim account alone does not provide a definition of assertion.

The truth-aim account of assertion is by definition non-factive. Rules and aims, as I characterised them, are mutually exclusive: if truth is the aim of assertion,

⁴ If the notion of “purported” seems too vague, it can be specified as an ‘agent-neutral’ (as opposed to ‘agent-relative’, cf. Nagel 1986) reason determined by the conventions of the relevant practice.

it cannot be its rule. The truth-aim account is thus compatible with *all and only* non-factive accounts: any non-factive norm (JR, BR, etc.) can be paired with it. For the purpose of this paper, I do not wish to commit to any specific pairing, but I would nonetheless like to make a suggestion. If assertions purport to aim at truth, it seems intuitive that asserting cooperatively requires one *at least to try to assert a true proposition* (cf. Dummett 1981:356); for this attempt to be rational, speakers will also have to follow JR (assert only what is reasonable for them to believe).

One of the key selling points of the truth-aim account is its ability to solve the puzzle introduced at the beginning of this paper (reconciling IMPROPER FALSITY with PERMISSIBLE FALSITY). A final issue to be considered is whether it represents the *only* solution to the proposed puzzle. As a matter of fact, some activist philosophers have argued (against the standard view) that PERMISSIBLE FALSITY is compatible with factive accounts (DeRose 2002, Weiner 2005, Whiting 2012). These authors typically appeal to a distinction between primary and secondary violations: they claim that if you *reasonably believe* that you are following a rule, you are ‘secondarily’ following that rule. In this sense, unlucky assertions are in secondary compliance with factive rules: even if an unlucky assertor violates factive rules in a ‘primary’ sense (she makes a false assertion), she follows them in a ‘secondary’ sense (she reasonably believes that the assertion is true). Unlucky assertions are thus ‘secondarily permissible’; if PERMISSIBLE FALSITY is qualified to allow for secondary compliance, it is consistent with factive accounts.

Rather than an objection to the truth-aim account, the primary-secondary distinction merely represents a competing explanation of the puzzle. That being said, the viability of this solution can be challenged. First, as already mentioned, to allow for secondary compliance PERMISSIBLE FALSITY has to be altered: the alternative view only offers a solution to a *revised* version of the puzzle. Acknowledging that unlucky assertions are *secondarily* permissible (*i.e.* primarily impermissible, but excusable) is not yet to admit that they are permissible primarily, *qua* assertions. The puzzle at stake, however, is motivated by an intuition of primary permissibility; unless one denies this intuition, this alternative explanation does not offer a solution to the puzzle. Second, the test for the rule/aim distinction has shown that rules typically do not allow for ‘secondary violations’ (violations by careful and competent agents). In characterising unlucky assertions as secondary violations, the alternative view entails that the norms regulating a common practice like asserting are exceptional, while failing to provide an explanation for their exceptionality. In sum, even though the primary/secondary distinction shows that an alternative explanation to the puzzle is conceivable, such explanation

presents serious difficulties⁵; insofar as it avoids such difficulties, the truth-aim account still represents the best available solution on the market.

Overall, my proposal to understand truth as the aim of assertion places itself within an orthodox tradition in speech act theory⁶, and offers a plausible solution to the puzzle presented in the opening of this paper. Unlike other non-factive accounts, it explains IMPROPER FALSITY: we criticise false assertions and deem them ‘incorrect’ not because they violate the rule of assertion, but because they do not ‘fit’ reality as they purport to. Unlike factive accounts, it is compatible with PERMISSIBLE FALSITY: unlucky assertions are (primarily) permissible, even if they do not meet their purported aim. In addition to this, unlike other accounts, the proposed view finds independent motivation from the TEST FOR THE RULE-AIM DISTINCTION.

References

- Alston, WJ. 2000. *Illocutionary Acts and Sentence Meaning*. Ithaca: Cornell University Press.
- Bach, K. 2010. Knowledge in and out of context. In J. K. Campbell & M. O’Rourke (Eds.), *Knowledge and skepticism*. Cambridge: MIT Press.
- Brandom, R. 1983. Asserting. *Nous* 17 (4): 637–650.
- Brandom, R. 1994. *Making it Explicit*. Cambridge: Harvard University Press.
- Cappelen, H. 2011. Against Assertion. In *Assertion : New Philosophical Essays*.
- DeRose, K. 2002. Assertion, Knowledge, and Context. *The Philosophical Review* 111(2): 167-203.
- Douven, I. 2006. Assertion, knowledge, and rational credibility. *Philosophical Review*, 115(4), 449–485.
- Dummett, M. 1981. Assertion. In *Frege: Philosophy of language*. London: Duckworth.

⁵ A number of additional arguments against the primary/secondary distinction can be found in the literature, and a growing consensus is emerging that the distinction is spurious. Lackey (2007) and Pagin (2016) offer the most extensive criticisms; compelling arguments are also found in Douven (2006:478-480), Stone (2007:100), Engel (2008) (Koethe 2009:631fn16), Cappelen (2011:46), Greenough (2011:fn29), Kvanvig (2011:242), Hinchman (2013:641fn6).

⁶ Some other philosophers have defended a truth-aim account of assertion: amongst them, the most notorious are Williams (1966:18-9) and Dummett (1981). Other have endorsed a comparable view, according to which *to assert a proposition is to present it as true* (Frege 1892:34, Wright 1992:23-34, Brandom 1983, 1994). Despite the apparent differences, this latter view is similar to the truth-aim account at least in the sense that it sets truth as the success-condition for the speaker’s act of asserting.

- Frege, G. 1892. Über sinn und bedeutung. *Zeitschrift Für Philosophie Und Philosophische Kritik*, 100:22–50.
- Green, MS. 2013. Assertions. In *Pragmatics of Speech Actions, Vol II*.
- Greenough, P. 2011. Truth-Relativism, Norm-Relativism, and Assertion. In *Assertion: New Philosophical Essays* (pp. 1–49).
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hinchman, E. S. 2013. Assertion, Sincerity, and Knowledge. *Noûs*, 47(4), 613–646.
- Hindriks, F. 2007. The status of the knowledge account of assertion. *Linguistics and Philosophy*, 30(3), 393–406.
- Humberstone, L. 1992. Direction of Fit. *Mind* 101 (401).
- Kemp, G. 2007. Assertion as a Practice. In D. Greimann and G. Siegart (eds.) *Truth and Speech-Acts*, 106–29. Routledge.
- Koethe, J. 2009. Knowledge and the Norms of Assertion. *Australasian Journal of Philosophy*, 87(4), 625–638.
- Kvanvig, J. 2009. Assertion, Knowledge, and Lotteries. In D. Pritchard & P. Greenough (Eds.), *Williamson on knowledge* (140–160). Oxford: Oxford University Press.
- Kvanvig, J. L. 2011. Norms of Assertion. In *Assertion: New Philosophical Essays*.
- Lackey, J. 2007. “Norms of Assertion.” *Nous*, 41(4): 594–626.
- Maitra, I. 2011. Assertion, norms, and games. In Jessica Brown & Herman Cappelen (eds.), *Assertion: New Philosophical Essays*, 277–296. Oxford: Oxford University Press.
- Nagel, T. 1986. *The View from Nowhere*. NY: Oxford University Press.
- Pagin, P. 2016. Problems with Norms of Assertion. *Philosophy and Phenomenological Research*, 93(1), 178–207.
- Récanati, F. 1987. *Meaning and Force: The Pragmatics of Performative Utterances*. Cambridge: Cambridge University Press.
- Schwyzler, H. 1969. Rules and Practices. *Philosophical Review* 78 (4):451.
- Searle, JR. 1969. *Speech Acts. An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Searle, JR. 1976. A classification of illocutionary acts. *Language in Society* 5:1–23.
- Searle, JR. 2007. Illocutionary acts and the concept of truth. In D. Greimann and G. Siegart (eds.) *Truth and Speech-Acts*, 31–40. Routledge.
- Stone, J. 2007. Contextualism and Warranted Assertion. *Pacific Philosophical Quarterly*, 88, 92–113.

- Vanderveken, D. 1990. *Meaning and Speech Acts* (Vol. 1). Cambridge University Press.
- van der Schaar, M. 2011. Assertion and grounding: a theory of assertion for constructive type theory. *Synthese*, 183(2), 187–210.
- Weiner, M. 2005. Must We Know What We Say? *Philosophical Review*, 114(2), 227–251.
- Whiting, D. 2012. Stick to the Facts: On the Norms of Assertion. *Erkenntnis* 78 (4):847–867.
- Williams, B. 1966, Consistency and realism. *Proceedings of the Aristotelian Society* 40:1-22.
- Williamson, T. 1996/2000. Knowing and Asserting. *The Philosophical Review* 105 (4):489–523; revised and reprinted in *Knowledge and Its Limits*. Oxford: Oxford University Press.