# Devious Stipulations

## John Horden

(forthcoming in *Oxford Studies in Metaphysics*; draft of July 2015)

## 1. Introduction

Traditionally, analytic truths have been thought, in some good sense, to require nothing of the world. And it is a corollary of this thought that combining a theory with its analytic consequences never produces a theory that requires more of the world than the original. Rudolf Carnap (1950) employed this idea in an attempt to reconcile his empiricism with a liberal acceptance of abstract entities; claiming that the existence of numbers, properties etc. can be trivially deduced by the rules governing our expressions for such things, once those expressions and rules are introduced into our language. In a highly influential response (1951), Willard Van Orman Quine argued that the analytic/synthetic distinction Carnap relied on for that result is untenable. Thereafter many so-called analytic philosophers, reputedly a majority for a while, followed Quine in disavowing analyticity altogether.

However, since then, analyticity has slowly crept back into philosophical respectability. Nowadays, it seems, analytic philosophers are mostly inclined to accept both that there are analytic truths (see Bourget and Chalmers, 2014), and that such interpreted sentences are characteristically metaphysically undemanding, at least in the sense of being neutral with respect to our position in logical space. Rather than rejecting analyticity outright, most would simply deny that there are any metaphysically *interesting* analytic truths. Meanwhile, going against the grain somewhat, there have in recent years been several notable attempts to answer ontological questions through conceptual analysis (see Hale and Wright, 2001; Schiffer, 2003; Thomasson, 2007, 2015; Hirsch, 2011; Hofweber and Velleman, 2011; Steinberg, 2013). And, by and large, all sides agree that *if* the existence of certain things analytically follows from sentences we already accept, then explicitly acknowledging the existence of those things cannot sensibly be regarded as an extra theoretical cost.

In his 'Analyticity and Ontology', Louis deRosset challenges this consensus. He tries to refute two principles (2015: 131, 139):

**DAO** If $P$ analytically entails the existence of certain things, then a theory that contains $P$ but does not claim that those things exist is no more ontologically parsimonious than a theory that also claims that they exist.

**GAO** If $P$ analytically entails $Q$, then ($P \wedge Q$) requires nothing more of the world than does $P$.

(DAO is what deRosset calls 'the doctrine of analyticity in ontology' and is his main target. GAO is a generalisation thereof.)

He glosses these principles as follows: 'a sentence $\phi$ is *analytic* iff it is entailed by true sentences $\psi_1$, $\psi_2$, ... such that failure to accept any $\psi_n$ constitutes some measure of linguistic incompetence' (2015: 133); 'a sentence $\phi$ *analytically entails* a sentence $\psi$ iff the material conditional ($\phi \Rightarrow \psi$) is analytic' (2015: 133); 'the parsimony of a theory is given by what the truth of the theory requires of the world *with respect to what there is*' (2015: 138). So the more ontologically parsimonious a theory is, the less it requires of the world with respect to what there is.

DeRosset does not define what it is for a sentence or theory to *require something of the world*, but makes clear that 'the world' here rigidly designates the actual world:

> The argument [against GAO] here relies on the assumption that *actually: $\phi$* requires no less of the world – the *actual* world, that is – than does $\phi$. [...W]hat's required of the actual world for *actually: grass is green* to be true is just for the actual world to meet whatever requirements there are for *grass is green* to be true. (2015: 143; see also n. 20)

Perhaps we should accept something along the following lines: a sentence $\phi$ requires of the world that $R$ iff (i) '$R$' does not logically follow from sentences that are true entirely because of their meanings, and (ii) whichever world is actual, if $\phi$ is true, this is partly because $R$. And then: a sentence $\phi$ requires nothing more of the world than a sentence $\psi$ iff, however '$R$' is replaced, if $\phi$ requires of the world that $R$, then $\psi$ requires of the world that $R$. Thus 'Actually grass is green' requires nothing more or less of the world than 'Grass is green'.

I won't rely on the correctness of the above definitions exactly as stated; just as deRosset, so he tells us (2015: 134), does not rely on his 'rough and ready' definition of analyticity. However, I shall make two assumptions. First: (schematically) if a sentence requires of the world that *R*, then, whichever world is actual, that sentence is true only if *R*. Second: if a sentence is true entirely because of its meaning, then it requires nothing of the world. (Notice that I say 'entirely because' here, not 'only because'. This is to allow for overdetermination. Plausibly, 'Everything is self-identical' is true partly because everything is self-identical (see Quine, 1954: 113; Boghossian, 1996: 364). Nonetheless, that sentence is fully determined to be true by its meaning; so in the relevant sense, it requires nothing of the world (see Russell, 2008: §1.2).)

Now, to see the *prima facie* appeal of DAO and GAO, consider:

(M1)  Someone is married.

(M2)  Someone is married to someone else.

(M1) explicitly affirms the existence of one person only, whereas (M2) explicitly affirms the existence of two. Nonetheless it seems clear that the conjunction of (M1) and (M2) requires nothing more of the world (and in particular, is no less ontologically parsimonious) than (M1) alone, because (M1) analytically entails (M2). And likewise, we might think, for more metaphysically interesting analytic entailments, if such there be.

However, deRosset denies that DAO and GAO hold in full generality. He suggests that in the metaphysically interesting cases just alluded to, these principles are liable to break down. So even if, to take his leading example, 'There are particles arranged tablewise in location *L*' analytically entails 'There is a table in *L*', it remains plausible that the conjunction of these sentences is less ontologically parsimonious than the former sentence alone. (Here he targets Amie Thomasson in particular. For the relevant sense of 'arranged tablewise' and the like, see her 2007: 16-17.)

Thus while unQuineanly granting the tenability of the analytic/synthetic distinction, deRosset disputes its metaphysical significance. To this end, he purports to produce counterexamples to DAO and GAO by means of linguistic stipulations. He offers two main candidate counterexamples just to GAO (though he admits that the first of these is not decisive) and one specifically to DAO. I aim to show where these arguments go wrong.

## 2. 'Verdantly*'

Here is deRosset's first candidate counterexample to GAO. (I bypass his initial, illustratively unsuccessful 'verdantly' stipulation.) Suppose we introduce a new sentential operator 'verdantly*' by stipulating the following rules of inference:

$P \vdash$ Verdantly* $P$

Verdantly* $P \vdash$ Actually grass is green

Verdantly* $P \vdash$ Possibly $P$

'Actually' here has its familiar interpretation from two-dimensional semantics: ⌜Actually $\phi$⌝ is true at a world iff $\phi$ is true at the actual world. Hence what deRosset calls *Stevenson's constraint* on linguistic stipulation (a generalised version of the constraint proposed by Stevenson, 1961) initially appears to be satisfied:

> [A] linguistic stipulation succeeds only if there is a consistent way to assign truth conditions to sentences containing the introduced term that makes the content of the stipulation true. (2015: 141)

To see this, suppose grass were not green. Then by the above introduction rule, 'Verdantly* grass is not green' would be true, and so by the first elimination rule, 'Actually grass is green' would be true. But even if grass were not green, grass would still be green at the actual world, so no inconsistency results.

If his 'verdantly*' stipulation succeeds, deRosset tells us, we have a counterexample to GAO. For then 'Snow is white' analytically entails 'Actually grass is green'. But 'Snow is white and actually grass is green' requires more of the world than 'Snow is white'. For the latter only requires that snow be white, whereas the former also requires that grass be green.

However, deRosset admits some uncertainty as to whether this stipulation succeeds (2015: 143-4). For perhaps when applying Stevenson's constraint we are not entitled to take for granted *a posteriori* necessities such as the fact that actually grass is green. Perhaps we

should also consider how to evaluate sentences containing 'verdantly*' on the assumption that actually grass is not green. But then consider what truth value should be assigned to

Verdantly* grass is either green or not green.

If it is true, the first elimination rule for 'verdantly*' is invalid, because 'Actually grass is green' is by hypothesis false. If, on the other hand, the evaluated sentence is false, the introduction rule is invalid, because 'Grass is either green or not green' is (we may assume) logically true. So the 'verdantly*' stipulation fails to meet Stevenson's constraint on this more stringent construal. DeRosset ultimately leaves open how stringently Stevenson's constraint should be construed, and so admits that his 'verdantly*' stipulation does not provide a decisive counterexample to GAO.

Fair enough; though whichever way we construe Stevenson's constraint, I wonder how useful it will be for adjudicating on attempted stipulations in cases where the bounds of conceptual possibility are in dispute. Should we consider how to evaluate sentences containing 'number' or 'proposition' on the assumption that there are only finitely many things, for example? If so, then the broadly Fregean and Carnapian views of abstracta targeted by deRosset are in trouble; but any such objection to those views – which, to be fair, deRosset does not himself raise – seems question-begging.

In any case, there is another reason to think that the 'verdantly*' stipulation fails. For 'verdantly*', as defined, is by reasonable standards clearly and fatally *nonconservative*. Nuel Belnap (1962) recommended a conservativeness constraint on stipulation in response to Arthur Prior's (1960) parody of implicit definition, wherein 'tonk' is defined by the rules:

$P \vdash P$ tonk $Q$

$P$ tonk $Q \vdash Q$

Thus any sentence analytically entails any other; hence every sentence is analytic. Obviously the stipulation fails; the only question is *how*. J. T. Stevenson (1961) proposed one constraint in response: every truth-functional connective must have a consistent truth table (notice that this is significantly weaker than the constraint deRosset names after Stevenson, on either construal of the latter). Belnap proposed another constraint: the introduction of any new vocabulary must yield a conservative extension of the language. That is to say, the rules

5

stipulated to govern new vocabulary cannot allow the derivation of any sentence of the old language (i.e. any sentence without the new vocabulary) that was not already derivable.[1] In footnotes, deRosset describes this conservativeness constraint, and observes that it is met by his subsequent stipulations – presumably regarding this as a point in their favour – but conspicuously does not claim the same for his 'verdantly*' stipulation. And it is easy to see why: 'Actually grass is green' is a synthetic sentence of English before this stipulation, but would become derivable from any sentence whatsoever (and hence analytic) were the stipulation to succeed. Thus Belnap's constraint is manifestly violated.

Admittedly, many regard Belnap's constraint as too strong (see Read, 1988: §9.3, 2000: 125-7; Peacocke, 1993: §§3-4, 2004: 18-21; Prawitz, 1994: 374; Shapiro, 1998: §3). For example, it seems legitimate to add second-order quantifiers or a truth predicate to a first-order system of arithmetic, though doing so enables the derivation of the original system's Gödel sentence. Also, abstraction principles such as Hume's Principle

The number of $F$s = the number of $G$s iff there are just as many $F$s as $G$s

if successfully stipulated, allow us to derive results about the (infinite) size of the domain of first-order quantification that were previously statable but perhaps unprovable. Hence Bob Hale and Crispin Wright endorse a weaker version of conservativeness (cf. Field, 1980: ch. 1; Wright, 1997: §9, 1999: §2.5; Schiffer, 2003: §2.2), also mentioned by deRosset (2015: 141, n. 19), with the following proviso attached:

It is our view that a stipulation *may* have consequences which can be expressed in the antecedent language, and to which there need have been no previous commitment, without compromise of its legitimacy *provided* the truth of these consequences makes no demands on the previously recognised ontology. (Hale and Wright, 2000: 302, n. 32)

However, the truth of 'Actually grass is green' clearly makes demands on the previously recognised ontology. For it demands that grass be green. So the 'verdantly*' stipulation fails to meet even this weaker version of conservativeness. So if either Belnap's or Hale and Wright's version of conservativeness is correct, this stipulation cannot provide a genuine counterexample to GAO.

---

[1] This principle of conservativeness, or 'noncreativity', was previously endorsed by others; see e.g. Frege, 1914: 208. For a historical discussion, see Urbaniak and Hämäri, 2012.

Moreover, regardless of the specific constraints on stipulation, deRosset's 'verdantly*' example faces a ruinous dilemma. First horn: 'Actually grass is green' retains its prior meaning. Then the nonanalyticity of this sentence, given its empirical status, is far more certain than the success of the stipulation. Given our understanding of analyticity, regrettably imprecise as that may be, it is absurd to think that we could make this sentence analytic without changing its meaning. So we should conclude, in Moorean fashion, that the stipulation *somehow* fails. If there were no known constraint that prohibited this stipulation, we would have to posit one. That, *mutatis mutandis*, is the lesson of 'tonk'. Second horn: 'Actually grass is green' loses its prior meaning. Then we don't know what, if anything, this sentence means after the stipulation: its subsequent meaning seems to be radically underdetermined. So we have no reason to accept that its conjunction with 'Snow is white' requires more of the world than 'Snow is white' alone. Either way, we do not get a genuine counterexample to GAO.

# 3. 'Grassgreen'

Let us now examine deRosset's second candidate counterexample to GAO:

> Suppose we stipulate that 'grassgreen' is to be a predicate that expresses the property *being green* if, as a matter of fact, grass is green, and *not being green* otherwise. (2015: 144)

This stipulation, deRosset tells us, meets Stevenson's constraint even on its more stringent construal, as well as meeting other proposed constraints including conservativeness, generality and harmony.[2]

Now, he continues, we have a more decisive counterexample to GAO. For if this stipulation succeeds, then 'Grass is grassgreen' is thereby guaranteed to be actually true, no

---

[2] The latter two constraints are proposed by Hale and Wright (2000) – in addition to their version of conservativeness – and are described by deRosset thus: '*Generality*: the stipulation should enable the interpretation of a wide enough range of relevant sentences'; '*Harmony*: the introduction and elimination rules should not allow us to infer more (or problematically less) than our warrant for the premises allows us to infer' (2015: 141, n. 19).

matter what colour grass actually has.[3] So 'Snow is white' analytically entails 'Grass is grassgreen'. But 'Snow is white and grass is grassgreen' requires more of the world than 'Snow is white'. For since actually grass is green, 'grassgreen', according to the stipulation, expresses the property of being green. So whereas 'Snow is white' only requires that snow be white, 'Snow is white and grass is grassgreen' also requires that grass be green.

Clearly something fishy is going on here. If 'Grass is grassgreen' is guaranteed to be actually true, no matter what colour grass actually has, then it cannot *require* of the actual world that grass be green. Even if we grant that the stipulation succeeds, it seems sensible to examine the semantics of the novel term 'grassgreen' a little more closely before we embrace inconsistency. So what is its stipulated meaning? That, I think, depends on whether the stipulation is *conditional* or *unconditional*. This yields another dilemma.

First horn: 'is grassgreen' is conditionally stipulated to mean *is green* if actually grass is green, and *is not green* otherwise. Then the stipulation does not by itself settle the meaning of this predicate. The inscription-type 'Grass is grassgreen' is thereby guaranteed to express some truth at the actual world in our extended language, but *which* truth it expresses depends on the actual colour of grass. In fact, since actually grass is green, the resulting interpreted sentence is exactly synonymous with 'Grass is green'. So we have no reason to think that it is analytically entailed by 'Snow is white'. Second horn: 'is grassgreen' is unconditionally stipulated to mean *is green iff actually grass is green*.[4] Then 'Grass is grassgreen' does not require of the actual world that grass be green. Nor does it require anything else of the world. So its conjunction with 'Snow is white' requires nothing more of the world than 'Snow is white' alone. Either way, we do not get a genuine counterexample to GAO.

DeRosset concedes that it is unclear what, exactly, we should take 'grassgreen' to mean, but insists that this does not matter for his purposes, so long as his stipulation succeeds (2015: 148). On the contrary, his argument here apparently rests on an equivocation between two intensionally equivalent interpretations of the introduced predicate, either of which is consistent with the explicit content of the stipulation. (Likewise for deRosset's third candidate counterexample, as we shall shortly see.) On one interpretation, 'Grass is grassgreen' requires something of the world, but is not guaranteed by its meaning to be actually true. On another interpretation, 'Grass is grassgreen' is guaranteed by its meaning to

---

[3] At least, 'Grass, if it exists, is grassgreen' is guaranteed to be actually true; like deRosset, I set aside this nuance. Also I grant that analytic truths can be contingent. As deRosset explains in his paper, he regards that assumption as inessential to his arguments anyway.

[4] Or, more longwindedly: *is such that either (a) it is green and actually grass is green or (b) it is not green and it is not the case that actually grass is green*.

be actually true, but requires nothing of the world. And yet, on an intensional individuation of properties, 'grassgreen' expresses the same property on either interpretation: the property of being green.

As it happens, deRosset anticipates the second interpretation, though he relegates this point to a brief footnote:

> It is plausible, perhaps, to think that the truth of ['Grass is grassgreen'] requires that grass be green iff grass is *actually* green; this is the view most naturally suggested by the content of the 'grassgreen' stipulation. But this requirement clearly goes beyond the requirement for the truth of ['Snow is white'], as evidenced by the fact that the requirement is not satisfied in circumstances in which grass is purple but snow is still white. (2015: 161-2, n. 47)

This response is clearly inadequate, however. For deRosset, recall, takes the expression 'requires nothing more of the world', as it appears in GAO, to be equivalent to 'requires nothing more of the actual world'. And though 'Grass is green iff actually grass is green' plausibly requires of *nonactual* worlds that grass be green, it requires nothing of the actual world (see Evans, 1979: §4; Davies and Humberstone, 1980: §2). So if 'Grass is grassgreen' merely abbreviates that biconditional, then its conjunction with 'Snow is white' requires nothing more of the world than 'Snow is white' alone. It is irrelevant that at some other possible world the biconditional is false while 'Snow is white' is true.

## 4. 'Priman Beings'

Now for deRosset's third candidate counterexample to GAO, which he also intends to serve as a counterexample to DAO:

> Suppose we stipulate that 'priman being' is to be a predicate that expresses the property *being a human being* if, as a matter of fact, there are more than seven billion human beings, and *being a prime number* otherwise. (2015: 146)

The stipulation here closely parallels that of deRosset's previous example. And as with his previous example, he assures us that this stipulation meets all the constraints mentioned so far.

If this stipulation succeeds, he continues, then the material conditional 'If there are more than 7 billion prime numbers, then there are more than 7 billion priman beings' is thereby guaranteed to be actually true, no matter how many humans there actually are, hence it is analytic. So 'There are more than 7 billion prime numbers' analytically entails 'There are more than 7 billion priman beings'. But 'There are more than 7 billion prime numbers and there are more than 7 billion priman beings' requires more of the world (and in particular, is less ontologically parsimonious) than 'There are more than 7 billion prime numbers'. For since actually there are more than 7 billion humans, 'priman being', according to the stipulation, expresses the property of being human. So whereas 'There are more than 7 billion prime numbers' does not require the existence of any humans, 'There are more than 7 billion prime numbers and there are more than 7 billion priman beings' requires the existence of more than 7 billion of them.

The analogy with deRosset's 'grassgreen' example should be clear; hence I shall respond analogously. If, given the existence of more than 7 billion prime numbers, 'There are more than 7 billion priman beings' is guaranteed to be actually true, no matter how many humans there actually are, then it cannot *require* of the actual world that there be more than 7 billion humans. Even if we grant that the stipulation succeeds, we should examine the semantics of the novel term 'priman being' a little more closely before we embrace inconsistency. So what is its stipulated meaning? Again, that depends on whether the stipulation is *conditional* or *unconditional*. This yields a dilemma.

First horn: 'is a priman being' is conditionally stipulated to mean *is human* if actually there are more than 7 billion humans, and *is a prime number* otherwise. Then the stipulation does not by itself settle the meaning of this predicate. Provided that there are more than 7 billion prime numbers, the inscription-type 'There are more than 7 billion priman beings' is guaranteed by the stipulation to express some truth at the actual world in our extended language, but *which* truth it expresses depends on how many humans there actually are. In fact, since actually there are more than 7 billion humans, the resulting interpreted sentence is exactly synonymous with 'There are more than 7 billion humans'. So we have no reason to think that it is analytically entailed by 'There are more than 7 billion prime numbers'. Second horn: 'is a priman being' is unconditionally stipulated to mean *is human if actually there are more than 7 billion humans, and a prime number otherwise*.[5] Then 'There are more than 7

---

[5] Or, more longwindedly: *is such that either (a) it is human and actually there are more than 7 billion humans or (b) it is a prime number and it is not the case that actually there are more than 7 billion humans*.

billion priman beings' requires nothing more of the actual world than whatever is required for there to be more than 7 billion prime numbers. So its conjunction with 'There are more than 7 billion prime numbers' requires nothing more of the world, and is no less ontologically parsimonious, than 'There are more than 7 billion prime numbers' alone. Either way, we do not get a genuine counterexample to either GAO or DAO.

# 5. Conclusion

So, as we have seen, none of deRosset's stipulations yields a genuine counterexample to either of the principles he targets. Despite his efforts to discredit them, the doctrine of analyticity in ontology and its generalisation remain as plausible as ever.[6]

# References

Belnap, N. D., 1962, 'Tonk, Plonk and Plink', *Analysis*, 22: 130-4.

Boghossian, P. A., 1996, 'Analyticity Reconsidered', *Noûs*, 30: 360-91.

Bourget, D. and D. J. Chalmers, 2014, 'What Do Philosophers Believe?', *Philosophical Studies*, 170: 465-500.

Carnap, R., 1950, 'Empiricism, Semantics, and Ontology', *Revue Internationale de Philosophie*, 4: 20-40; reprinted in his 1956, *Meaning and Necessity: A Study in Semantics and Modal Logic*, 2nd edn, Chicago, IL: University of Chicago Press.

Davies, M. and L. Humberstone, 1980, 'Two Notions of Necessity', *Philosophical Studies*, 38: 1-30.

deRosset, L., 2015, 'Analyticity and Ontology', in K. Bennett and D. W. Zimmerman, eds, *Oxford Studies in Metaphysics*, vol. 9, Oxford: Oxford University Press.

Evans, G., 1979, 'Reference and Contingency', *The Monist*, 62: 178-213.

Field, H. H., 1980, *Science Without Numbers: A Defence of Nominalism*, Princeton, NJ: Princeton University Press.

Frege, G., 1914, 'Logic in Mathematics', trans. P. Long and R. White, in H. Hermes, F. Kambartel and F. Kaulbach, eds, 1979, *Gottlob Frege: Posthumous Writings*, Oxford: Blackwell.

Hale, B. and C. Wright, 2000, 'Implicit Definition and the A Priori', in P. Boghossian and C. Peacocke, eds, *New Essays on the A Priori*, Oxford: Oxford University Press; reprinted in Hale and Wright, 2001.

—, 2001, *The Reason's Proper Study: Essays towards a Neo-Fregean Philosophy of Mathematics*, Oxford: Oxford University Press.

Hirsch, E., 2011, *Quantifier Variance and Realism: Essays in Metaontology*, Oxford: Oxford University Press.

Hofweber, T. and J. D. Velleman, 2011, 'How to Endure', *Philosophical Quarterly*, 61: 37-57.

Peacocke, C., 1993, 'Proof and Truth', in J. Haldane and C. Wright, eds, *Reality, Representation, and Projection*, Oxford: Oxford University Press.

—, 2004, *The Realm of Reason*, Oxford: Oxford University Press.

Prawitz, D., 1994, review of M. Dummett, *The Logical Basis of Metaphysics*, *Mind*, 103: 373-6.

Prior, A. N., 1960, 'The Runabout Inference-Ticket', *Analysis*, 21: 38-9.

Quine, W. V., 1951, 'Two Dogmas of Empiricism', *Philosophical Review*, 60: 20-43.

—, 1954, 'Carnap and Logical Truth', in his 1976, *The Ways of Paradox and Other Essays*, 2nd edn, Cambridge, MA: Harvard University Press.

Read, S., 1988, *Relevant Logic: A Philosophical Examination of Inference*, Oxford: Blackwell.

—, 2000, 'Harmony and Autonomy in Classical Logic', *Journal of Philosophical Logic*, 29: 123-54.

Russell, G., 2008, *Truth in Virtue of Meaning: A Defence of the Analytic/Synthetic Distinction*, Oxford: Oxford University Press.

Schiffer, S., 2003, *The Things We Mean*, Oxford: Oxford University Press.

Shapiro, S., 1998, 'Induction and Indefinite Extensibility: The Gödel Sentence is True, but Did Someone Change the Subject?', *Mind*, 107: 597-624.

Steinberg, A., 2013, 'Pleonastic Possible Worlds', *Philosophical Studies*, 164: 767-89.

Stevenson, J. T., 1961, 'Roundabout the Runabout Inference-Ticket', *Analysis*, 21: 124-8.

Thomasson, A. L., 2007, *Ordinary Objects*, Oxford: Oxford University Press.

—, 2015, *Ontology Made Easy*, Oxford: Oxford University Press.

Urbaniak, R. and K. S. Hämäri, 2012, 'Busting a Myth about Leśniewski and Definitions', *History and Philosophy of Logic*, 33: 159-89.

Wright, C., 1997, 'On the Philosophical Significance of Frege's Theorem', in R. Heck, ed., *Language, Thought, and Logic: Essays in Honour of Michael Dummett*, Oxford: Oxford University Press; reprinted in Hale and Wright, 2001.

—, 1999, 'Is Hume's Principle Analytic?', *Notre Dame Journal of Formal Logic*, 40: 6-30; reprinted in Hale and Wright, 2001.