

*I am delighted if you decide to quote this paper—please send me an email if you do.
Comments are welcome. Yours, ZZs.*

*To appear in the conference proceedings for the Pécs Philosophy Conference Series,
The Philosophy of Robert Brandom, April 2005.*

DEONTIC DISCOURSE—LESSONS FROM INFERENCEALISM*

Zsófia Zvolenszky
zvolenszky@nyu.edu

“I want to introduce here a way of thinking about semantics that is different from more familiar ones, and on that basis also a new way of thinking about logic.” This is how Robert Brandom (2000, p. 45) begins his essay “Semantic Inferentialism and Logical Expressivism”, in which he advocates the move towards *semantic inferentialism* (the view that a statement’s representational features are read off from the inference patterns it exhibits) and away from *semantic representationalism* (the view that a statement’s inference patterns are read off from its representational features). I adopt Brandom’s opening lines for two reasons. First, because my purpose in the present paper is to do for *deontic logic* (a special kind of modal logic to capture permissions and obligations) exactly what Brandom promises: introduce a new way of thinking about deontic logic and its semantics. Second, because deontic logic thereby provides an illuminating case study and testing ground for how a representationalist bias can yield a mistaken conception that we can rectify if we heed Brandom’s inferentialist lesson.

After (i) a brief introduction to deontic logic, I will (ii) identify a prevalent assumption, namely that extant deontic logics purport to be *Humean frameworks*: they aim to analyze and represent modal goings-on in terms of non-modal goings-on. I will then (iii) locate a recalcitrant problem that these logics face: *any statement of the form ‘If p then it must be that p’ turns out to be a truth of deontic logic*, even though some instances of this schema are clearly false (let *p* stand for any historical injustice, say, ‘Suleiman’s army pillaged Pécs in 1526’). It turns out that (iv) this problem calls for a radical solution: adopting a non-Humean framework, in which modal goings-on are represented as such and are irreducible to non-modal goings on. This might seem like a disappointing outcome, but I argue that it is not: (v) a non-Humean deontic logic says all that has to be said about deontic discourse; to think otherwise and to hope for a Humean framework had been a representationalist’s mirage.

Modal discourse is about necessity and possibility: what must/should/ought to/might/can be the case.

(1) Mary can drink Riesling Italic.

(1) is at least two ways ambiguous—one reading concerns Mary’s ability to drink wine, the other is about what Mary is allowed or permitted to do. Our concern here is

* I am grateful to Kit Fine for comments and discussion. This research was supported by the MTA–ELTE Philosophy of Language Research Group.

with the latter, deontic reading, and more generally, with deontic discourse (about patterns of reasoning concerning permissions and obligations). Deontic readings themselves are context-sensitive: in one context of utterance, (1) invokes religious requirements (to the effect that Mary's religion is compatible with her drinking Riesling Italico), in another, doctor's orders (to the effect that Mary's doctor does not think that her health and medications contraindicate her drinking Riesling Italico). The choice of deontic background is wide open.

Throughout the paper, let us assume a specific deontic background befitting the Pécs region—that of a wine etiquette. Imagine that it includes clauses about how various types of wine are served (temperature, glasses, use of a decanter), which wine goes with which type of dish, etc. Two of our sample clauses are:

- (2) You must drink wine.
- (3) If you eat artichokes, you must not drink wine.

(3) is a case of *conditional obligation*—an obligation to avoid wine is conditional upon artichoke consumption. (Most oenophiles agree that artichokes and wine cannot be happily paired.) Otherwise, the *absolute obligation* in (2) is operative—we have to make sure that we do our share of wine drinking (this clause makes for an uncommonly hearty wine guide).

As a species of modal discourse, logicians have analyzed deontic discourse in terms of quantification over possible worlds. Accordingly, (2) is equivalent to (2') or (2''):

- (2') In every one of the deontically accessible possible worlds, wine drinking takes place.
- (2'') In every one of the *best among the* deontically accessible possible worlds, wine drinking takes place.

(2'') is more complicated than (2') in that it requires an ordering over possible worlds based on how well they satisfy the contextually relevant obligations, in our case, about wine etiquette. Notice a crucial assumption present in both options: *they strive to reduce modal features of possible worlds (what the worlds allow and require) to non-modal or mundane features of a selected set of possible worlds* (the deontically superior ones, say). Call a framework *Humean* just in case it assumes the specification of the modal in terms of the mundane. For example, an obligation to drink wine is specified in terms of wine drinking habits prevalent in the deontically best worlds. An alternative, non-Humean approach maintains the following: possible worlds are specified in part by their mundane features concerning drinking and eating habits, and in part by their modal features about what laws, requirements and permissions are in effect in them. (2') and (2'') are alike in that both are Humean analyses of (2).

We are now in a position to witness a profound problem (to be referred to as 'The Problem') that has been plaguing deontic logic from its inception:

- (4) If you eat artichokes, then you must eat artichokes.
- (5) If you eat artichokes at 4 a.m., then you must be awake at 4 a.m.

Given the deontic background about wine etiquette against which (4) and (5) are uttered, they say, respectively, that the wine etiquette imposes an artichoke-eating requirement whenever artichoke eating occurs, and imposes a requirement about being

awake in the middle of the night whenever one happens to be awake then, eating artichokes. But it is obvious that just because one happens to be eating artichokes at a certain hour does not make for a wine etiquette requirement for her to do so. More generally, we expect the schema 'if p then it must be that p ' to be sometimes false (witness (4)). And even more generally, we expect 'if p then it must be that q ', where p entails q to be sometimes false (witness (5)). The Problem: instances of these schemas, including (4) and (5) turn out to be truths of deontic logic.

Here is how The Problem arises. Assume an analysis of 'must' along the lines of (2"); my diagnosis carries over to other Humean analyses of 'must':

- (6) 'it must be that p ' is true if and only if
among the worlds under consideration, p is true in every one of the deontically closest worlds.

We thus have a way of handling absolute obligations like (2); we still need to provide a definition to cover conditional obligations; let us therefore add:

- (7) 'if p then it must be that q ' is true if and only if
among the worlds under consideration, all the p -worlds make 'it must be that q ' true.

These definitions suppress quite a bit of technical detail spelled out elsewhere (see Zvolenszky 2002).

We can analyze (3) based on (6) and (7):

- (3) If you eat artichokes, you must not drink wine.

First, apply (7). Begin with a set of possible worlds (with various eating, drinking and sleeping habits). Among these worlds, zoom in on those in which artichokes are eaten; 'you must not drink wine' is true in them. This is where definition (6) comes in: 'you must not drink wine' is true just in case with respect to each world still under consideration (each artichoke-eating world, that is), one refrains from drinking wine in the deontically best worlds.

A crucial question signifies that this analysis is not yet complete: out of which worlds do we select the deontically best options? When evaluating the consequent, are we still restricting our attention to artichoke eatings or not? In other words, do we retain the restriction about artichoke eating imposed by the antecedent (the first half) of the conditional in (3), or do we look more broadly and check the wine drinking in worlds of all stripes, with or without artichoke eating? The two options yield radically different outcomes for (3):

- (A) If we *retain the restriction in the antecedent*, then the truth of (3) requires only that there be no wine drinking in the worlds in which the wine etiquette is obeyed as much as possible *and* artichoke eating has transpired. This seems like a plausible way of analyzing a conditional obligation like (3).
- (B) If *the restriction in the antecedent is discarded* once we are evaluating the consequent (the second half) of (3), then the truth of (3) would hinge on a universal prohibition on wine drinking in all deontically superior worlds, not just those in which artichokes have been consumed. This seems wrong-headed. Intuitively, (3) does not forbid wine drinking if, say, one is eating smoked

salmon. And the deontically superior worlds could well include worlds where smoked salmon is washed down with Riesling, in accordance with a plausible wine etiquette clause to the effect that

(8) If you eat smoked salmon, you ought to drink Riesling.

Given (8), worlds that combine smoked salmon and Riesling are better than those in which smoked salmon is unaccompanied by wine. Option (B) would fail to capture this. (8), as well as the absolute obligation to drink wine (in (2)) are supposed to be consistent with a conditional obligation like (3). But they would come out as contradictory if we opted for discarding the antecedent of (3) when assessing its consequent.

It therefore looks like (3) calls for option (A) as opposed to (B). But (A) turns out to be the wrong choice for (4), repeated below:

(4) If you eat artichokes, then you must eat artichokes.

Definition (7) instructs us to restrict our attention to artichoke-eating worlds; (A) then instructs us to retain that restriction for the worlds with respect to which we evaluate ‘you must eat artichokes’. This means that *all the worlds* we are considering have artichoke eating. But then *all the deontically best worlds* are guaranteed to have artichoke eating. And by definition (6), ‘you must eat artichokes’ cannot but be true, affording (4) the status of a logical truth. To avoid this result, a straightforward choice would have been to broaden our coverage of worlds along the lines of option (B), to include worlds in which no artichokes are eaten. But we have just seen that (B) is a non-starter for garden variety conditional obligations like (3) and (8), so we cannot justify adopting it.

A convenient solution suggests itself at this juncture: why not retain both options (A) and (B) and apply each as needed? Indeed, we see precisely this idea buried in Jackson’s (1985) proposal—a rare instance within the extensive deontic logic literature to accord The Problem due weight. This is a workable option, but one that results in a non-Humean framework—one that no longer specifies modal features like requirements about, say, wine drinking in terms of mundane features about, say, wine consumption.

It is instructive to see why a non-Humean framework is inevitable (for a more detailed explanation, see Zvolenszky 2005). Consider again (3) (‘If you eat artichokes, you must not drink wine’) and (4). Rewind to the step where we have already zoomed in on the artichoke-eating worlds and for each of them, we are at the point of applying definition (6), to analyze the consequent of each conditional—one about avoiding wine, the other about artichoke eating. The crucial difference between the two examples is that given the wine etiquette, we expect the consequent ‘you must not drink wine’ to hold in each artichoke eating world; but we also want to avoid making ‘you must eat artichokes’ true in the artichoke-eating worlds. To formulate this very pair of criteria that would allow us to shuttle between options (A) and (B), we need to make independent appeal to the status of the deontic statements in question: whether or not they are supposed to hold. Is it a deontic feature of a world in which someone eats artichokes that she should not drink wine, given the wine etiquette? (The answer is ‘yes’, so we should select option (A) here.) Is it a deontic feature of a world in which someone eats artichokes that she should eat artichokes, given the wine

etiquette? (The answer is ‘no’, so we should pick option (B) here.) In short, *our justification for which of two accounts of conditional obligation apply in a particular case relies on independently available information about deontic features*. And such information requires that we have deontic features of worlds—about what is required and what is allowed—to be represented as such: as obligations and permissions. We can no longer simply rely on definition (6) to read off deontic features from the mundane goings-on of possible worlds. Instead, we need the deontic features specified independently: artichoke-eating worlds that satisfy the wine etiquette have, besides the mundane feature that artichokes are being consumed, also deontic features like ‘one can refrain from artichoke eating’, and ‘one must not drink wine’. Deontic features specified as such, specified independently of definition (6) signal that we are already treading on non-Humean ground.

We have just witnessed a solution to The Problem that comes at the price of giving up on a Humean framework. Logicians and semanticists would find this too hefty a price. Considering that, it seems like a glaring case of negligence that The Problem has received hardly any attention, especially if we think of all the attention that has been lavished on deontic paradoxes like the Good Samaritan and the Paradox of Gentle Murder. To the extent that The Problem was acknowledged, some proposed to bite the bullet and admit that (4) is a logical truth, albeit in a degenerate sense (van Fraassen 1972, Feldman 1988). Others, including Jackson (1985), heralded accounts that supply the desired truth conditions for (3) and (4) (Frank 1997, Geurts 2004). On the one hand, I take the former bullet to be unbitable—how can we claim that we have given an adequate analysis of deontic discourse if our account is saddled with the unintuitive truth value assignments with which (4) and (5) as logical truths leave us? On the other hand, Frank’s account turns out to generate comparably serious problems elsewhere (Zvolenszky 2002), whereas Geurts’ proposal, insofar as it works, amounts to abandoning a Humean framework for deontic logic in much the same way as Jackson’s account does (Zvolenszky 2005).

In the remainder of this paper, I will give some general reasons for why The Problem calls for a non-Humean framework. Philosophers and linguists working on deontic modality would be unhappy with going non-Humean, because that would deem one of their central projects a failure: for over fifty years, they have been trying to analyze ‘must’ when in fact it cannot be analyzed. It is at this point that Brandom’s insights about semantic inferentialism and representationalism reveal quite a bit. First, we see whence the appeal of a Humean framework: it is all important for a representationalist semantics. Second, by thinking in terms of an inferentialist semantics instead, as Brandom recommends, it becomes much easier to resist the attraction of a Humean framework in the face of a difficulty as basic and profound as The Problem.

For all I have said about Jackson’s account turning out to be a non-Humean one, some might still remain hopeful about the prospects of a Humean account, wondering about the various details that could be tweaked. We could supply an alternative treatment of conditionals—analyze them as strict or material implication. Or we could change the way in which the conditional interacts with the deontic elements—by having the deontic part take large scope. Or in place of the traditional one-place ‘must’, we could introduce an alternative operator which takes two arguments, one of which is the conditional’s antecedent, the other its consequent (for example, van Fraassen 1972). This array of options creates the impression that there is plenty of room for maneuvering our way around The Problem. In fact, The Problem arises independently

of the specifics of the modal logic framework I have sketched, and therefore remains effective even if we change some of those specifics.

Bringing about The Problem arises as soon as we posit two basic ingredients. One is definition (6), which embodies the Humean expectation that we cash out the modal (in our case, the deontic) in terms of the mundane: ‘it must be that p ’ is true just in case all of the (deontically best) accessible worlds make p true. The other ingredient concerns our treatment of deontic conditionals like (3): regardless of the details of our account, we cannot avoid restricting our attention to the worlds that satisfy the antecedent—artichoke-eating worlds, that is. Here is why. There is one very attractive link that suggests itself between absolute obligations and conditional ones: *conditional obligations are just like absolute obligations in worlds in which their antecedents obtain*. This is supposed to be at the foundation of the semantics for conditional obligation.

We should not confuse this semantic expectation with the syntactic rule (or derivation) of Detachment, which parallels the semantics: for worlds in which our premises include some conditional obligation (say (3)) and its antecedent (that you eat artichokes), we can derive an absolute obligation involving the consequent (that wine drinking should be avoided). Jackson claims that “[d]etachment is plausible. If it ought to be that Attila goes to jail given that he has raped and pillaged, and he has raped and pillaged, then it ought to be that Attila goes to jail.” (Jackson 1985, p. 191) Tomberlin (1989, p. 110) concurs: such inferences have to be valid “...for statements of conditional obligation to play a genuine role in the normative guidance of conduct”.

Still, it is important to distinguish Detachment from the semantic expectation about linking conditional obligations to absolute ones. The reason: the latter is even more general. For example, in van Fraassen’s (1972) framework, Detachment cannot be derived (see Tomberlin 1989); nonetheless, he adheres to the semantic expectation:

There is a rudimentary semantic criterion that yields another axiom and rule. The intuitive meaning of ‘given A ’ is such that, if a sentence ends with it, then *any possibility that does not satisfy A is irrelevant to the evaluation of that sentence*. Thus, the evaluation of $O(A/B)$ [It ought to be that A given B] cannot depend on $H(A)$ [the set of attainable states in which A holds] as such but at most on $H(A \ \& \ B)$ [the set of attainable states in which both A and B hold]. (van Fraassen 1972, p. 421; emphasis added)

The semantic expectation linking conditional and absolute obligations, together with the assumption of a Humean framework suffices to give rise to The Problem. Whichever way we cash out the syntax and semantics of conditional obligations, we can follow the reasoning already familiar from our previous considerations about options (A) and (B): in order to get the right truth conditions for garden variety conditional obligations like (3), we need to evaluate the consequent with respect to worlds that satisfy the antecedent (this is established by the semantic connection between absolute and conditional obligations). But then if we add definition (6) so all it takes to make ‘it must be that p ’ true is that p holds in every world under consideration, then (4) and (5) come out as logical truths, and we have The Problem.

We therefore have only three options, two of which are unpalatable: biting the bullet with respect to The Problem, and severing the semantic link between conditional and absolute obligations. With this pair of options out, we are left with just one choice: giving up on a Humean framework for deontic logic. We have already seen the outlines of what this entails: ‘you must not drink wine’ is no longer represented in terms of universal abstinence from wine drinking in the deontically best among the worlds being considered. That is to say, definition (6) is no longer available. And giving up on that

definition amounts to relinquishing all hope for a possible-worlds *analysis* of deontic modality.

Initially, this might seem like a dispiriting loss. But Brandom's semantic inferentialism is an important source of consolation. His goal is

... to define representational properties in terms of inferential ones, which must accordingly be capable of being understood antecedently. [We] start with a notion of content as determining what is a reason for what and understand truth and representation as features of ideas that are not only manifested in but actually consist in their role in reasoning (Brandom 2000, p. 47; emphasis in original).

With respect to deontic discourse, this approach calls for capturing the *inference patterns exhibited in deontic discourse*. For example, that 'it must be that q ' follows from p , together with 'if p then it must be that q '; that *not- p* is compatible with 'it must be that p '; that 'if p , then *not- p* is permitted' is true for some instances of p . It is through these inference patterns that we arrive at what 'it must be that p ' means and how its meaning is different from but related to that of the statement that p . In particular, our understanding of the former does not hinge on cashing it out in part in terms of the latter (as definition (6) does). But that does not mean that we have left something important unexplained.

When I first encountered The Problem back in 2001 and began to realize that the history of deontic modality consisted in large part in attempts to patch up a Humean framework in ways that were contrived, ad hoc, and ultimately inadequate, I wondered why logicians and semanticists were not drawn to a non-Humean alternative. Brandom has an explanation for that: from the prevalent representationalist perspective, whose master concepts are representation, truth, and reference (as opposed to inference), it seems obvious that part of what we need to do is locate within the representation of 'it must be that p ' the representation of p itself. Otherwise, we have left something unexplained. Hence comes the insistence on definition (6) and on maintaining a possible-worlds analysis of modality. This paper has provided two reasons to relent and go non-Humean: the Problem leaves us no other choice; in addition, the non-Humean alternative seems in no way a disheartening or inferior option once we take in Brandom's lesson from inferentialism.

REFERENCES

- Brandom, Robert. 2000. *Articulating Reasons. An Introduction to Inferentialism*. Harvard University Press, Cambridge, Massachusetts.
- Feldman, Fred. 1990. A Simpler Solution to the Paradoxes of Deontic Logic. *Philosophical Perspectives*, 4, Action Theory and Philosophy of Mind, 309–341.
- van Fraassen, Bas. 1972. The Logic of Conditional Obligation. *Journal of Philosophical Logic* 1, 417–438.
- Frank, Anette. 1997. *Context Dependence in Modal Constructions*. PhD thesis, Universität Stuttgart.
- Geurts, Bart. 2004. On an Ambiguity of Quantified Conditionals. Manuscript, University of Nijmegen.
- Jackson, Frank. 1985. On the Semantics and Logic of Obligation. *Mind* 94, 177–195.
- Tomberlin, James E. 1989. Deontic Paradox and Conditional Obligation. *Philosophy*

- and Phenomenological Research*, 50, 107–114.
- Zvolenszky, Zsófia. 2002. Is a Possible-worlds Semantics of Modality Possible?
Semantics and Linguistic Theory XII. Cornell University Press, Ithaca, New York.
- Zvolenszky, Zsófia. 2005. Modality Unanalyzed. Manuscript, New York University.