

THE DEONTOLOGY OF INFERENCE

STEPHEN FAGAN, RAY JENNINGS AND MARTIN ALLEN

ABSTRACT. Various methods have been adopted by various researchers for generating well-behaved sublogics of classical propositional logic, with the notion of good-behaviour varying with the proposed applications or interpretations of the systems studied. Paraconsistentists have generally sought systems that restrict in some way the inferences available from classical inconsistencies of one sort or another, either by altering the truth conditions of connectives, or by altering the preservational requirements imposed on inference. This paper proposes a third approach, one which incidentally generates sublogics of classical logic, and neither discards classical truth-conditions nor substitutes more stringent preservational requirements.

Taking classical logic as providing untransgressable boundaries of inference (that is, as playing the role of physical possibility in the *ought implies can* constraint) we consider inference as a practice and consider the consequences for that practice of the adoption of axiologically based deontic principles. In effect we explore the possibility of a deontic logic of inference based upon comparisons of consistency measures, and the conceptions of inference that emerge from such axiological principles. Among the conclusions of the research is that such an approach casts useful light on connexivist inference.

1. INTRODUCTION

Formal methods have mercifully departed from philosophical in their treatment of inferrability, largely replacing the language of inferrability with that of provability (\vdash) and semantic entailment (\models). In the present study we depart a little from purely formal considerations without, we hope, lapsing into the familiar philosophical practice of talking in order to find out what we are talking about.

Even in speaking of the rules of a formal system (say of natural deduction) comfortably as *rules of inference*, we are aware that these are really rules for extending a proof, and that if *inference* enters, it consists in inscribing a formula as the next formula in a sequence of formulae, perhaps with some documentation. What goes on in the head of the proof-constructor is irrelevant to the status of the line, provided that there is a rule of the system that licenses it. Thus far, proof-theory treats what it refers to as inference as an act: the rules considered severally permit some such acts and do not permit others. But no rule either forbids or requires any act. And the rules even taken collectively do not enjoin any act at all. They do not require anyone to engage in constructing a proof. In Kantian language, even taken collectively they represent only a hypothetical imperative: if you are to construct an L -proof then every line that you add must be licensed by a rule of L . But again, in Kantian language they don't really constitute a deontological imperative, since the requirement that every line be L -justified is definitional of an L -proof. A Kantian map of this territory will identify that collective constraint is the indirect source

of *can*'s, not of *ought*'s. Infact there are no *ought*'s, or if they enter, they do so through instructional labs. 'Look', we say, 'those seventeen \vee -introductions have evidently got you nowhere closer to having a proof. Don't you think you ought to vary your approach. You ought to have assumed this' and so on. The rules of the system collectively act to constrain whatever hypothetical *ought*'s the activity of constructing an efficient proof imposes. If no rule justifies the inscription of a particular line, then it cannot be that the proof-constructor *ought* to inscribe that line, because *ought* implies *can*. On the analogy with the Kantian deontology of ordinary life, the rules taken collectively seem to play the role of a set of physical laws or conceptual constraints.

By contrast, the semantical basis for the rules does constrain the formulation of the rules, since, by the usual standard, every rule must be sound. (We set aside the complications of rules taking subproofs as their inputs.) Whatever counts as the goodness of inferential outcomes enters the mix in the selection of rules. If (proof-theoretic) inferential agents are axiologists, then they are rule-axiologists rather than act-axiologists. They don't trouble themselves, say, with the preservation of truth at every line that they inscribe. They trust that that worry has been taken out of their hands by the formulators of the rules that license their inscribed lines.

2. REAL DEONTOLOGY

Without wishing to name names, we register the impression that among philosophical logicians there are Kantians who suppose (a) that the rules of natural deductive logic, when we have got them right, will be those that constrain proof in much the same way as physical law constrains oughts, and (b) that the evidence for the correctness of that system of rules (or perhaps equivalence class of such systems) is available in the world, much as is the evidence for a correct physical theory. Imagine a theory that posits propositions as fundamental objects of the world, and the inferential nature of the world as consisting in the algebraic structure of that set of objects, as revealed by experimental inquiry. The question as to what can be inferred from what, is the question, 'What lies above what in the algebraic structure of the propositions of the world?' Such a theorist might suppose, for example, that whenever a proposition $\alpha \wedge \neg\alpha$ is found in this structure, both α and $\neg\alpha$ lie somewhere above it, but that there are propositions that do not.

There are others who see the matter differently, perhaps because their approach to deontology more generally is more pragmatic. The contrast is this: the Kantian supposes that categorical imperatives can be discovered as brute facts of the human condition, and that these constitute absolute *ought*'s; our non-Kantian supposes that human nature and other empirically investigatable considerations provide the data for experimental formulations of public policy. On this latter, non-Kantian, view the questions are not questions of what is obligatory, but of what, given commonly agreed ends, we ought to require of agents; not questions of what is permissible, but questions of what, given commonly agreed ends, we ought to permit agents to do. In fact, in Kantian language, these requirements do not even count as *hypothetical* imperatives, since sanctions are, as a rule, legislatively associated with their omission, and for practical purposes they are allied rather with physical laws

in restricting what it can be claimed that we ought to do. Though from time to time we feel compelled to insist that we ought to act other than in conformity with legislated requirements, we typically do so because we suppose the requirement ought not to have been enacted. In general, however, there is no such conflict and we do not claim that we ought act in a certain way if acting in that way is illegal, as a claim that we ought to act in a certain way is defeated by the observation that acting in that way is physically or conceptually impossible.

Put bluntly, when the two approaches are assessed along these lines, the contrast that emerges is this. On the one approach we secure correct inference by an investigation intended to disclose what *real* restrictions apply to what *can* be inferred. This is achieved by replacing classical logic, or better, by replacing the classical account of the algebraic structure of the propositions, by another account. At any rate, those restrictions upon what *can* be inferred impose restrictions upon what *ought* to be. On the second approach, either no such real algebraic structure is posited, perhaps no real propositions either, or the classical structure is held to play this role. In any case, no additional restrictions are imposed upon what *ought* to be inferred by restrictions upon what it is supposed *can* be inferred. What *can* be inferred is given by whatever inferences classical logic admits. This imposes restrictions upon what *ought* to be inferred, but the latter deontic category is admitted to be a proper subcategory of the former. In fact the relationship between the two parallels the ordinary deontological relationship between what is physically and what is morally possible. The whole responsibility for any additional restrictions to the inferential deontology, on this views, falls upon those charged, in particular circumstances, with the formulation of inferential policies. Circumstances and ends determine inferential axiology; axiology determines deontology.

3. WHENCE AXIOLOGY?

Here we want consider the second, non-Kantian approach. That is, we take the structure of the propositions of the world to be the boolean algebra of classical logic, and that algebraic structure to constrain the deontology of inference as the physical nature of the world constrains the deontology of action more generally. That is, we take it as a constraint upon what any inferential axiology can dictate that if we ought to infer α from Σ , then α classically follows from Σ .

3.1 The *-able* of *inferr-able*

Now it is unclear from the habits of ordinary speech whether the *-able* suffix of *inferrable* is taken to be like the *-ible* of *permissible* or like the *-able* of *payable* or *reportable* as in *This bill is payable within thirty days*, or *Wilful damage to public property is a reportable offence*, that is, whether the deontic category to which it assigns a contemplated inference is that of the permissible or that of the obligatory. That is one dimension of difficulty, but there is a second. This is the question whether the *-able* of *inferrable* on whichever of those dual readings is to be understood by analogy with alethic modalities or with deontic ones. There may be intellectual restrictions on what can *really* be inferred. If there are, they are difficult to discern in the ordinary dealings of daily life, both professional and social. (“How

could he possibly have inferred that,” we ask, but we know that in most spheres anything more or less well-formed is possible.) So definitions of both categories fall back upon such restrictors as *for a sane/ for a rational/ for a reasonable*. Each such restrictor will yield a distinct if not clear and distinct modality. Certainly some of them will look more alethic, and others more deontic in character. So for example, if there really are intellectual limits to inference, then the category of what it is possible for a living person to infer will be restrictive, but less restrictive and so more alethic than the category of what it is reasonable to expect a sane, rational, well-trained, skeptic to infer. What it is impossible for such folk not to infer is equally or more problematic, particularly bearing in mind that in speaking of inference, we have no very deep understanding of what we are talking about.

We must somehow or other set the difficulties aside, because if we are to apply some sort of deontological language to inference at all, then we must also be prepared to apply some sort of ontological language as well. Deontology is bounded by ontology. If there were such a thing as axiomatic ethics, one of its axioms would have to be *There is at least one impossible act*. Else we could just make ethics what we wanted it to be whatever that was. But more particularly, our conception of ethics is such that what we take someone to be required to do, we must also take to be possible, perhaps even contingent. Perhaps too, what we take ourselves to be permitted to do, we must also take to be possible or contingent.

3.2 Internal or external negation

In axiological comparisons, what do we compare: the outcome of inferring α with the outcome of not inferring α , or the outcome of inferring α with the outcome of inferring $\neg\alpha$. Something like this problem arises in the consideration of more general ethical questions, even those that have historically had purely deontological motivations, and been supposed to have pure deontological status. Consider the question whether we ought to treat other humans as ends rather than as means. As a matter of fact, even after we have selected our whole training wardrobe, our diet, and our mode of transport, arenas of action through which we might stand in need of the Kantian franking, the over whelming part of humanity remains as persons whom I treat as neither means nor ends, having nary an occasion to do so. It may even be that at that level of generality, even Kant could be moved by considerations of efficiency to admit that I in default of no real duty to treat those people as ends. This surely is an occasion for the invocation of some *Ought implies can* principle. On the other hand in the question whether, as regards a particular person for whom the opportunity has arisen, I should treat that person as an end or as a means, the relevant comparison is the deliberation of outcomes in which end and non-end are compared, rather than outcomes in which treating and non-treating are. All obligations, even those arising from categorical imperatives are conditional in character. It is perhaps a debatable point, but we take ourselves not to be too far out of line in distinguishing the cases at least. Efficiency of outcomes: brevity, elegance, parsimony in resource-management, enter into external comparisons. What enters into internal ones?

3.3 Inferential axiology

Peter Schotch has entered as a plausible inferential slogan, the Hippocratic principle, *Primum non nocere: first, do no harm*. The intent is to lend gravity to an earlier preservationist dictum that in making inferences, particularly inferences from damaged or fragile data, we should not make matters worse than they already are. Rather we should exercise a kind of archaeological care, preserving data from deterioration or contamination from the addition to it of conclusions that do not respect such virtues as the data may possess. Schotch had in mind such measures as that of incoherence level: $\iota(\Sigma)$

the cardinality of the least decomposition of Σ into classically coherent subsets, if such a decomposition exists; else ∞ .

The idea generalizes easily to yield a measure for any naturally non-monotonic property of sets of sentences, that is, properties not preserved by the operation of forming supersets. But other measures have since been independently introduced, such as the Thorn measure:

$\iota(\varphi, \Gamma)$, the cardinality of the least partition of Γ such that the union of the (classical) inferential closure of its cells has the property φ

and the Brown measure for (Γ, α) , the level of the set $\{\alpha \wedge \gamma \mid \gamma \in \Gamma\}$

and that of φ -dilution:

$\delta(\varphi, \Gamma)$, the cardinality of the least subset of Γ having the property φ if such a subset exists; else ∞ .

The original motives for such innovations were, as we have said, preservationist ones. The idea was that inference should be selectively curatorial, designed around the requirements of a particular measure, preserving the value of the measure through the inferential closure of a set of data. Thus, for example, coherence-level-preserving inference does not treat data from n different sources as it would treat data from $n - 1$ sources. If each of n data source has the property φ , then the closure of the set of data under φ -level-preserving inference will still be decomposable into n cells each of which has φ . Dilution-preserving inference enables a system of restricted capacity to use such a subset of classical inferential methods as will guarantee that inconsistencies in its data remain undetected. If the least incoherent subset of a set Σ of data is too large for a system to discover the inconsistency of Σ , then the inconsistency of the inferential closure of Σ will likewise be undiscoverable by the system.

Our present interest in such measures is more directly axiological. Since these are measures of goodness of sets of sentences, they can also figure in the axiological calculations concerning outcomes of inferences. That is, if Σ is a set, and α wff, we can compare, with respect to some measure, the inference of α from Σ with the inference of $\neg\alpha$ from Σ , as the comparison of $\Sigma \cup \{\alpha\}$ and $\Sigma \cup \{\alpha\}$ with respect to the measure, and infer according to our preferences for that measure. The question that we ask, though we can as yet give it no satisfactory answer is this: Can measures previously introduced as candidates for inferential preservation be given a more general axiological redeployment in the theory of inference?

In utilitarian ethical theory, no detailed examination has been given to the corresponding questions. Yet they seem to be formulable in corresponding terms. Consider the two dicta:

- I Act in such as a way that you make the world no worse than it is.
- II If performing an act will make the world better than refraining from the act, then perform the act.

The former is preservationist in character; the latter more generally axiological. In a predicament in which one is forced to choose among a range of possible actions, the former would enjoin that we select only those acts that do not make matters worse. But in casting about among the courses of action open to us, we must ask, for each contemplated action, ‘Shall I perform it or not?’ Thus conformity with the first principle seems to require us to invoke the second, since for each available course, we are forced to choose between the performance of an act and its negation. In the restricted domain of inference, and for well-defined inferential axiologies, the question as to the relationship between preservation and axiology would seem more readily available for exploration.

3.4 The ontology of inference

Since something must be done, we settle more or less peremptorily, the issue of what, in the deontology of inference, will play the role of ontological constraint. For the purposes of this study, which is purely formal, and purely propositional, we simply give the job to classical logic, giving classical inferrability the role in inferential deontology that possibility plays in the deontology of action. What could play the role of necessity remains unclear, but having done this much, we can formulate our *ought implies can* principle: Nothing ought to be inferred that is not classically inferrable. This would correspond to a consistency principle

$$\neg\Box\perp$$

in a deontic system. We also as a principle that *may implies can*, which yields a principle akin to a normality principle

$$\neg\Diamond\perp$$

In the more general domain of deontic logic, this is thought by some to be deontically objectionable, though their aims in rejecting it do not always match the effects of doing so. However, in the inferential restriction, its inclusion is not vicious or untoward. It’s effect is only this: that our study remains a study within the domain of classical inferrability. The question is which classically permitted inferences an adopted deontology of inference would dictate, and which not.

3.5 Choosing a measure

As in the more general deontological setting, Goldilocks principles apply. Standards must not be so high as to make inference extremely difficult and therefore rare and therefore practically useless, nor so feeble as to make it automatic and therefore banal and therefore practically useless. So for example, we will not choose coherence-level as the basis for the comparison of outcomes. For if $\mathcal{I}(\Sigma \cup \{\alpha\}) > \mathcal{I}(\Sigma \cup \{\neg\alpha\})$, then $c \cup \{\alpha\}$ must be inconsistent, for every cell c , including every unit cell, in every π in $\Pi_{\mathcal{I}(\Sigma)}(\Sigma)$ in which every cell is classically consistent. Not that this

makes a bad case for inferring $\neg\alpha$; rather it requires too good a case except when Σ is classically consistent. Even for the preservation of level, this would represent a very conservative strategy. At the other extreme, to adopt classical inferrability itself as the standard would be to adopt a principle akin to that of making every possible act an obligatory one.

Now we have embarked upon this exploration, not because we had a standard in mind, nor because we supposed that some single standard would be correct for every domain of inferential activity. Why would inference be different in this respect from the rest of life? In determining what is to count as inferential goodness our interests and purposes have to be given a role. Here our purposes are exploratory rather than strictly preservationist, but our interests, and therefore the principles that guide our selection of axiological standards are nevertheless paraconsistent. So it is worth reporting that any strict comparison-based deontology of inference of the sort we have outlined is paraconsistent. That is, even from inconsistent data, there must be inferences that such comparisons will not enjoin. The reason is that for any ensemble of sentences Σ (even an inconsistent ensemble), and any sentence α , the inference of α will be enjoined only on the basis of an asymmetry between $\Sigma \cup \{\alpha\}$ and $\Sigma \cup \{\neg\alpha\}$. That is, if the comparison enjoins the inference of α , then it will not enjoin the inference of $\neg\alpha$, and vice versa. Thus the inferential *ought* will satisfy a principle akin to the so-called deontic law [D]:

$$\Box\alpha \rightarrow \neg\Box\neg\alpha$$

Such a deontology will enjoin no conflicting inferential *oughts*. Equivalently, it does not enjoin what it does not also permit.

3.6 The axiology of dilution

As an illustration, we consider the axiology of coherence-dilution: that is, that of the measure $\delta(\varphi, \Gamma)$, the cardinality of the least subset of Γ having the property φ if such a subset exists; else ∞ , letting φ be classical coherence. For simplicity, we write $\delta(\Gamma)$ for $\delta(\varphi, \Gamma)$ in this case, and speak of the dilution of Γ . What would the character of a deontology of inference be that enjoined the inference of α from Γ when $\Gamma \vdash \alpha$ and $\delta(\Gamma \cup \{\neg\alpha\}) < \delta(\Gamma \cup \{\alpha\})$?

First, for the purely structural reasons given already, δ -comparison will not dictate the inference of α and the inference of $\neg\alpha$ from the same ensemble of sentences.

Second, δ -comparison does not dictate the inference of α from $\{\neg\alpha\}$ or $\neg\alpha$ from $\{\alpha\}$.

Thus, the inferential dictates of δ -comparisons are inferentially connectivist.

Third, with respect to enjoined inferences, δ -comparison renders sets containing explicit contradictions inert, since $\delta(\{\perp, \alpha\}) = \delta(\{\perp, \neg\alpha\})$ for every wff α . But for the same reason, δ -comparison *permits* any inference whatever from any ensemble containing an explicit contradiction.

Fourth, δ -comparisons enjoin all \wedge -introductions that do not decrease dilution.

3.7 Observations and Queries

δ -comparisons enjoin inferences that are not δ -preserving.¹ Consider the set

$$\Sigma = \{\alpha, \alpha \rightarrow \beta, \beta \rightarrow \gamma, \gamma \rightarrow \eta, \neg\eta\}$$

Then $\delta(\Sigma) = 5$. But $\delta(\Sigma \cup \{\neg\alpha \wedge \eta\}) = 3$, and $\delta(\Sigma \cup \{\alpha \vee \neg\eta\}) = 4$. Therefore the δ -comparison enjoins the inference of $\alpha \wedge \neg\eta$. In fact, in such inconsistent implicational chains having an odd number of \rightarrow 's, δ -comparisons will enjoin inferences of consequents or negations of consequents, accordingly as the wff has occurrences late or early in the chain. Thus, for the cited set Σ , $\delta(\Sigma \cup \{\beta\}) = \delta(\Sigma \cup \{\neg\beta\}) + 1$. But $\delta(\Sigma \cup \neg\gamma) = \delta(\Sigma \cup \gamma) + 1$. In the case of consistent implicational chains with even numbers of \rightarrow 's, there will be only one consequent to which δ -comparisons will be inferentially indifferent: the consequent equidistant from the beginning and the end of the chain. So if the axiology of δ -comparisons is paraconsistent, it is somewhat eccentrically so. To be sure, there may be applications for which such discriminations are just what the witch-doctor ordered. But on the face of it, δ -comparison seems to be inadequate as an inferential frank when classical inferrability alone is taken as the ontological constraint on inference.

3.8 Conclusion

How are the two axiological principles I and II related? At least in the particular inferential application we have considered they are not in the relationship of a goal and a successful strategy for its achievement. Rather they would seem to be in the relationship of inferential policy and a particular application. The difference is that in the second case the policy itself cannot be lost sight of: only δ -preserving inferences can be permitted; only sentences whose negations are impermissible are obligatory. The principle I must be revised to read

I' Act *only* in such a way that you make the world no worse than it is.

In the particular inferential application, the revision has the effect (assuming the liberal principle that what is not forbidden is permitted) that α is a permissible inference from Σ iff $\delta(\Sigma \cup \{\alpha\}) \geq \delta(\Sigma)$. By the duality of *ought* and *may*, α is an obligatory inference from iff $\delta(\Sigma \cup \{\neg\alpha\}) < \delta(\Sigma)$. Thus if α is an inference that ought to be made, then $\delta(\Sigma \cup \{\neg\alpha\}) < \delta(\Sigma \cup \alpha)$, but that inequality is by itself an insufficient condition.

Laboratory for Logic and Experimental Philosophy,
Department of Philosophy,
Simon Fraser University,
Burnaby, British Columbia, V5A 1S6,
Canada

¹We are grateful to Martin Allen, who drew this class of examples to our attention.