Logica universalis 1 (2007), 295–310
1661-8297/020295-16, DOI 10.1007/s11787-007-0015-x

Logica Universalis

On Preserving

Gillman Payette and Peter K. Schotch

Abstract. This paper examines the underpinnings of the *preservationist* approach to characterizing inference relations. Starting with a critique of the 'truth-preservation' semantic paradigm, we discuss the merits of characterizing an inference relation in terms of preserving *consistency*. Finally we turn our attention to the generalization of consistency introduced in the early work of Jennings and Schotch, namely the concept of *level*.

Mathematics Subject Classification (2000). Primary 03B53; Secondary 03B22.

Keywords. Preservationism, paraconsistency.

1. Introduction

The (classical) semantic paradigm for correct inference is often given the name 'truth-preservation.' This is typically spelled out to the awe-struck students in some such way as:

> An inference from a set of premises, Γ, to a conclusion, α, is correct, say *valid*, if and only if whenever all the members of Γ are true, then so is α.

This understanding of the slogan may be tried, but is it actually true? There is a problem: the way that 'truth' is used in connection with the premises is distinct from the way that it is used with the conclusion. In other words, this could be no better than a quick and dirty gloss. The chief virtue of the formulation is that of seeming correct to the naive and untutored.

But what of the sophisticates? They might well ask for the precise sense in which truth is supposed to be preserved in this way of unpacking. On the right hand side of the 'whenever' we are talking about the truth of a single formula while on the left hand side we are talking about the truth of a bunch of single formulas. Is it the truth of the whole gang which is 'preserved?'

The authors wish to acknowledge the support of SSHRCC under research grant 410-2005-1088.

Of course it is open to the dyed-in-the-wool classicalist to reply scornfully that we need only replace the set on the left with the conjunction of its members. In this way truth is preserved from single formula to single formula as homogenously as anyone could wish.

It is open, but not particularly inviting. In the first place, this strategy forces us to restrict the underlying language to one which has conjunction and conditional connectives – which must operate in something like the usual (which is to say classical) way. There are enough who would chafe under this restriction, that a sensitive theorist would hesitate to impose it.

We are inclined to think of this business of coding up the valid inferences in terms of their 'corresponding conditionals'[1] as an accident of the classical way of thinking and that it is no part of the *definition* of a correct account of inference. We also notice that on the proposal, we are restricted to finite sets.

However, it is possible, as is done in quantum logic, to forgo the conditional and restrict discussion to conjunctions on the left. In doing so the definition of 'follows from' for sets is defined via sentence-sentence relations.[2] One could presumably, then, say $\Gamma \vdash_X \alpha$ when there was a particular conjunction to 'do the work'. Apart from the requirement of conjunction, this suggestion removes some intuitive distinctions that will be explained in the sequel. Namely, the idea that one may break up a collection of sentences. Further, could one really claim that a logic, so defined, be monotonic? Infinite sets seem to be ruled out. Although our discussion is restricted to compact logics, much generality would be sacrificed if one's goal is to describe a uniform notion of 'follows from.'

Setting aside this unpalatable proposal then, we ask how is this notion of truth-preservation supposed to work? Since there is no gang on the right we seem to be talking about a different kind of truth, individual truth maybe, from the kind we are talking about on the left – mass truth perhaps. Looked at in that somewhat jaundiced way, there isn't any preserving going on at all, but rather a sort of transmuting.

The classical paradigm *really* ought to be given by the slogan 'truth transmutation.' In passing from the gaggle of premises to the conclusion, gaggle-truth is transmuted into single formula truth. It may be more correct to say that, but it makes the whole paradigm somewhat less forceful or even less appealing.

What we need in order to rescue the very idea of preservationism is to talk entirely about sets. So we shall have to replace the arbitrary conclusion α with the entire *set* of conclusions which might correctly be drawn from Γ. We even have an attractive name for that set – the *theory* generated by Γ or the *deductive closure* of Γ. In formal terms this is the theory generated by Γ or the X-closure of Γ

$$\mathbb{C}_X(\Gamma) = \{\alpha | \Gamma \vdash_X \alpha\}$$

[1] We think this usage was coined first by Quine.

[2] The authors thank the anonymous referee for this suggestion.

Now that we have sets,[3] can we say what it is that gets preserved – can we characterize classical inference, for instance, as that relation between sets of formulas and their closures such that the property Φ is preserved?

We can see that gaggle-truth would seem to work here in the sense that whenever Γ is gaggle-true so must be $\mathbb{C}_{\vdash}(\Gamma)$, for $\vdash$ the classical notion of inference at least. We are unable to rid ourselves, however, of the notion that gaggle-truth is somewhat lacking from an intuitive perspective. Put simply, our notion of truth is carried by a predicate which applies to *sentences*, or formulas if we are in that mood. These are objects which might indeed belong to sets, but they aren't themselves sets. So however we construe the idea of a true set of sentences (or formulas) that construal will involve a stipulation, or more charitably, a new definition.

Generations of logic students may have been browbeaten into accepting: 'A set of sentences is true if and only if each member of the set is true.' But, it *is* a stipulation, and is no part of the definition of 'true.' It doesn't take very much imagination to think that somebody might actually balk at the stipulation. Somebody who is attracted to the idea of *coherence* for instance, might well want to say that truth must be defined for (certain kinds of) sets first, and that the sentential notion is derived from the set notion and not conversely. All of which is simply to say that a stipulation as to how we should understand the phrase 'true set of sentences' is unlikely to be beyond the bounds of controversy.[4] It may gladden our hearts to hear then, that there is another property, perhaps a more natural one, which will do what we want. That alternative property is *consistency.*

2. Making a few things precise

By a *logic* X, over a language $\mathcal{L}$ we understand the set of pairs $\langle \Gamma, \alpha \rangle$ such that Γ is a set of formulas from the language $\mathcal{L}$ and α is a formula from that same language, and $\Gamma \vdash_X \alpha$. In the sequel we frequently avoid mention of the language which underlies a given logic, when no confusion will thereby be engendered.

This set of pairs is also referred to as the *provability* or *inference relation* of X. In saying this we expose our extensional viewpoint according to which there is nothing to a logic over and above its inference relation. This has the immediate consequence that we shall take two logics X and Y which have the same inference relation, to be the same logic.

---

[3]This approach is different from the sequent version of Set-Set consequence as seen in [2]. Discussion of that can be found, for the current notions in [6].

[4]It may be helpful here to consider an analogy between sentences and numbers, taken to be *urelementen*. We can define the idea of a prime number easily enough but be puzzled about how to define a prime *set* of numbers. Somebody might be moved to offer: "Why not simply define a prime set of numbers to be a set of prime numbers?" The answer is likely to be: "Why bother?" indeed the whole idea of a prime set of numbers seems bizarre and unhelpful. We can easily imagine circumstances in which we would require a set of prime numbers but the reverse is true when it comes to a prime set.

When X is a logic, we refer to the X-deductive closure of the set Γ by means of '$\mathbb{C}_X(\Gamma)$'.

Unless the contrary is specified, every logic mentioned below will be *compact*, which is say that whenever proves $\alpha(\Gamma \vdash \alpha)$ it follows that there must be some finite subset, say Δ, of Γ, which proves α.

In mentioning consistency, we have in mind some previously given notion of inference, say $\vdash_X$. Each inference relation spawns a notion of consistency according to the formula

> Γ is consistent, *in* or *relative to* a logic X (alternatively, Γ is X-consistent) if and only if there is at least one formula α such that $\Gamma \nvdash_X \alpha$.

To say this in terms of provability rather than non-provability we might issue the definition:

> Γ is *inconsistent* in a logic X if and only if $\mathbb{C}_X(\Gamma) = \mathbb{S}$, where $\mathbb{S}$ is the set of all formulas of the underlying language of X, i.e., Γ proves everything.
>
> Where X is a logic, the associated consistency predicate (of sets of formulas) for X, is indicated by $\textsc{con}_X$.

We were interested in how an inference relation might be characterized in terms of preserving some property of sets. We have singled out consistency as a natural property of sets, and having done that we can see that preservation of consistency comes very naturally indeed. The time has come to say a little more exactly what we mean by 'characterized.' In order to do this we shall be making reference to the following three *structural* rules of inference.

- Reflexivity: $\alpha \in \Gamma \implies \Gamma \vdash \alpha$, referred to by [R].
- Transitivity: $\Gamma, \alpha \vdash \beta$ and $\Gamma \vdash \alpha \implies \Gamma \vdash \beta$, referred to by [Cut].
- Monotonicity $\Gamma \vdash \alpha \implies \Gamma \cup \Delta \vdash \alpha$ (Δ an arbitrary set of formulas), referred to by [Mon].

Unless there is a specific disavowal every inference relation we consider will be assumed to admit these three rules. It should be noted that on account of [Mon], if the empty set $\varnothing$ is inconsistent in X, then the inference relation for that logic contains every pair $\langle \Gamma, \alpha \rangle$. In such a case we say that X is the *trivial* logic over its underlying language. We shall take the logics we mention from now on to be non-trivial, barring a disclaimer to the contrary.

Let us say that an inference relation, say '$\vdash_X$', *preserves consistency* if and only if:

> If Γ is X-consistent (in the sense of the previous definition), then so is $\mathbb{C}_X(\Gamma)$.

It is easy to see that every inference relation with [Cut] and [Mon] must preserve consistency since if the closure of a set, Γ proves some formula, α, then by compactness some finite sequence of [Cut] operations will lead to the conclusion that Γ proves α. It may be that we end up showing that some subset of Γ proves α, which is why we require [Mon] in this case.

We say that X preserves consistency *in the strong sense* when the condition given above as necessary, is also sufficient.

It is similarly easy to see that since every set is contained in its deductive closure by [R], and since inconsistency is preserved by supersets, given [Mon], every inference relation satisfying the three structural rules preserves consistency in the strong sense.

This is all very well, but we haven't really gotten to anything that would single out an inference relation from the throng which preserve consistency. In order to do that it will be necessary to talk about a logic X preserving the consistency predicate of a logic Y, in the strong sense.

A moment's thought will show us that when the preservation is mutual – X and Y preserve each other's consistency predicates, (which implies that they share a common underlying language) then they must agree on which sets are consistent and which are inconsistent. For consider, if $\textsc{con}_X(\Gamma)$ and Y preserves the X consistency predicate then $\textsc{con}_X(\mathbb{C}_Y(\Gamma))$. Suppose that Γ is not Y-consistent, then $\mathbb{C}_Y(\Gamma) = \mathbb{S}$. By [R] $\mathbb{C}_X(\mathbb{C}_Y(\Gamma)) = \mathbb{C}_X(\mathbb{S}) = \mathbb{S}$ which is to say that $\mathbb{C}_Y(\Gamma)$ is not X-consistent, a contradiction. Similarly for the argument that Γ is consistent in Y and X preserves the Y consistency predicate.[5]

When two logics agree in this way, i.e., agree on the consistent and inconsistent sets, we shall say that they are *at evens*. Another moment's thought reveals that two logics which are at evens will preserve each other's consistency predicates. Assume X and Y are at evens and $\textsc{con}_X(\Gamma)$, but $\overline{\textsc{con}_X}(\mathbb{C}_Y(\Gamma))$ where the overline indicates predicate negation. Then $\overline{\textsc{con}_Y}(\mathbb{C}_Y(\Gamma))$ because X and Y agree on inconsistent sets. Whence by idempotentcy of '$\mathbb{C}_Y$' Γ is not consistent in Y, a contradiction. Thus we have,

Proposition 2.1. *Any two logics X and Y over a language $\mathcal{L}$ are at evens if and only if X and Y preserve each other's consistency predicates.*

This is nearly enough to guarantee that X and Y are the same logic. All we need is a kind of generalized negation principle:

Definition 2.2. A logic X is said to have *denial* provided that for every formula α, there is some formula β such that $\overline{\textsc{con}_X}(\{\alpha, \beta\})$.

[5]The anonymous referee suggested: "The definition of consistency preservation can (at least often) be strengthened. We can demand that every X-consistency-preserving extension of Γ also be a consistency-preserving extension of $\mathbb{C}_X(\Gamma)$." We assume the referee means something like the following – say for intuitionistic logic (IL). Let $\Gamma = \{A\}$ then $\Gamma \nvdash_{IL} B \vee \neg B$, but $\textsc{con}_{IL}(\Gamma')$ where $\Gamma' = \{A, B \vee \neg B\}$. In this case $\mathbb{C}_{IL}(\Gamma')$ will be a consistent extension $\mathbb{C}_{IL}(\Gamma)$, and so also of Γ. The referee continues: "That is, when we close Γ under a consistency-preserving consequence relation, we not only preserve the consistency of Γ, we also beg no questions, that is, we do not rule out as inconsistent extensions of $\mathbb{C}_X(\Gamma)$ that are consistent extensions of Γ." It is not our intention to rule out such consistent extensions, and the definition does not rule out such extensions. It is because of that possibility that the relation of 'at evens' is not equality of logics; a logic requires more structure to show that 'at evens' implies equality. Again we thank the referee for the comment.

In such a case we shall say that α and β deny each other (in X, which qualification we normally omit when it is clear from the context). We will assume that any logic we mention has denial.

Clearly if a logic has classical-like negation rules then it has denial since the negation of a formula will always be inconsistent with the formula. Of course classically, there are countably many other formulas which are inconsistent with any given formula, e.g., all those which are self-inconsistent and conjunctions which include $\neg\alpha$. The generalized notion doesn't require that there be distinct[6] denials for each formula, only that there be some or other formula which is not consistent with the given formula.

Evidently, if two logics are at evens, then if one has denial, so does the other. In fact something stronger holds, namely:

Proposition 2.3. *If two logics X and Y are at evens, and X has denial then, for every formula α there is some formula β for which both $\overline{\text{CON}_X}(\{\alpha,\beta\})$ and $\overline{\text{CON}_Y}(\{\alpha,\beta\})$.*

A further distinction to be made is that of a 'contingent' formula. Contingent formulas are those that are neither absurdities nor theorems. An absurd formula is one which is self-inconsistent. A theorem is a consequence of everything. Suppose that φ is a contingent formula, and there is another contingent formula ψ such that the pair set is inconsistent. These formulas will be called 'contingent denials'. The next definition is more complex.

Definition 2.4. Let α be a contingent formula of a logic X and β be a denial of α. β is a negation-denial of α (ND) if and only if,

1. β is contingent,
2. $\overline{\text{CON}_X}(\alpha,\beta)$ and,
3. if $\overline{\text{CON}_X}(\alpha,\delta)$ then $\delta \vdash_X \beta$.

These negation-denials are so called because they 'act like' negations in the way Koslow characterizes negations in [3]. Notice that ND is not a symmetric relationship. For instance the negation of intuitionistic logic satisfies this formulation, but A is not an ND of $\neg A$ in IL. That is because in intuitionistic logic $A \vdash \neg\neg A$, but the implication does not go in the other direction.

In the sequel the logics will obey the following. Suppose that $\overline{\text{CON}_X}(\Gamma,\beta)$. Then there is a denial δ of β such that $\Gamma \vdash_X \delta$. If one further assumes that ND is a symmetric relationship then the logic X will have the following property called [Den].

$$[\text{Den}]\quad \forall\Gamma,\ \alpha,\ \beta\ \Big[\beta ND\alpha \implies \big(\Gamma \vdash_X \alpha \iff \overline{\text{CON}_X}(\Gamma,\beta)\big)\Big]$$

Note that ND represents the relation of negation-denial.

Proposition 2.5. *Given that a logic X has the following properties:*

[6]Distinct up to logical equivalence, it goes without saying.

1. *If* $\overline{\text{CON}_X}(\Gamma, \beta)$ *then there is* δ *such that* $\overline{\text{CON}_X}(\delta, \beta)$ *and* $\Gamma \vdash_X \delta$.
2. X *has symmetric ND's.*

Then the logic satisfies [Den].

Proof. Suppose that ND is symmetric for X, and the other assumptions 1 and 2 hold of X. Then assume $\beta ND \alpha$. Suppose that $\Gamma \vdash_X \alpha$. Then, of course, $\overline{\text{CON}_X}(\Gamma, \beta)$ since β is a denial of α. Now suppose that $\overline{\text{CON}_X}(\Gamma, \beta)$. Then there is some denial of β, γ, such that $\Gamma \vdash_X \gamma$ by assumption. Since ND is symmetric, $\alpha ND \beta$, so by definition of an ND it follows that $\gamma \vdash_X \alpha$ and by [Cut] $\Gamma \vdash_X \alpha$. □

It is worth mentioning that if two logics are at evens, and one satisfies [Den] then it is not always the case that the other will. This is the situation between classical and intuitionistic logic respectively. The lesson is that [Den] may only hold for some formulas in a logic, but not all. It is a stronger assumption to assume that [Den] holds for all contingent formulas. Still, there are many logics which satisfy [Den]. It is clear from the context below that the denials used are of the kind just described.

Now we are ready to state our result:

Theorem 2.6 (Generalized Consistency Theorem). *Let* X *and* Y *have symmetric negation denials.* X *and* Y *are at evens if and only if,* X *and* Y *are the same logic.*

Proof. For this argument we split the equivalence into its necessary and sufficient halves.

($\Longrightarrow$) Assume $\mathbb{C}_X(\Gamma) = \mathbb{C}_Y(\Gamma)$ for every set Γ – which is to say that $X = Y$. To say that $\overline{\text{CON}_X}(\Gamma)$ is to say that the X-closure of Γ is $\mathbb{S}$. But then so must be the Y-closure of Γ. Similarly, to say that $\text{CON}_X(\Gamma)$ is to say that there is some α which is not in the X-closure of Γ, but then neither can α be in the Y-closure of Γ, hence $\text{CON}_Y(\Gamma)$. So X and Y are at evens.

($\Longleftarrow$) Suppose then that X and Y are at evens. Let Γ be a consistent set, which means by the assumption, that it is consistent in both logics. Assume for reductio that $\Gamma \vdash_X \alpha$ and $\Gamma \nvdash_Y \alpha$, and let β deny α. Thus, by [Den] $\overline{\text{CON}_X}(\Gamma \cup \{\beta\})$ and $\text{CON}_Y(\Gamma \cup \{\beta\})$, a contradiction. □

3. What's wrong with this picture?

To answer the question in the section heading, there really isn't anything wrong with an approach which characterizes inference in terms of preserving consistency. It's consistency itself, or at least many accounts of it, which casts a shadow over our everyday logical doings.

The way we have set things up, a set Γ of formulas is either consistent in a logic X, or it isn't. But it doesn't take much thought to see that such an all-or-nothing approach tramples some intuitive distinctions. In particular, we may find the *reason* for the inconsistency to be of interest.

In the logic X, for example, there may be a single formula δ which is, so to speak, inconsistent *by itself*. In other words $\overline{\text{CON}_X}(\{\delta\})$. Formulas of this dire sort are described above as being *self-inconsistent* (in X) or *absurd* in X. By [Mon] any set of formulas which contains a self-inconsistent formula is bound to be inconsistent.

We are now struck by the contrast between X-inconsistent sets which contain X-absurdities and those which do not. Isn't there an important distinction between these two cases? If we think of consistency as a desirable property which we are willing to trouble ourselves to achieve, then the trouble will be light indeed if all we need do is reject absurdities. On the other hand, an entire lifetime of angst may await those who wish to render consistent their beliefs or their obligations.[7] There is a great deal more that one could say on this topic and some of the current authors have said much of it. For now, we shall take it that the need for a distinction has been established, and our job is to construct an account of consistency which allows it.

We have in mind building upon what we have already discovered instead of pursuing a slash-and-burn policy. This means, among other things, that the predecessor account should appear as a special (or limiting) case of the new proposal. The intuitive distinction bruited above, is clearly a distinction between different kinds of inconsistency, or perhaps different degrees. We might think of one kind being *worse* than the other, which leads to a rather natural way of classifying inconsistency.

4. Speak of the level

The account of inconsistency which we propose is a generalization of the one first suggested in the 20th Century, in the work of Jennings and Schotch[8], namely the idea of a *level* (of incoherence, or inconsistency). The basic idea is that we can provide an intuitive measure of how inconsistent a set is by seeing how finely it must be divided before all of the divisions are consistent. What, in the earlier account is stated in terms of classical provability, we now state in terms of arbitrary inference relations which satisfy the minimal conditions given in the earlier section.

It all begins with the notion of a certain kind of indexed collection of sets being a *logical cover* for a set Γ of formulas, in the logic X. First we need a special kind of indexed family of sets (of formulas).

Definition 4.1. $A(\Delta) = \{a_0, a_1, \ldots, a_\xi\}$ is an indexed set *starting with* Δ, provided $a_0 = \Delta$ and all the indices $0 \ldots \xi$ are drawn from some index set I.

Definition 4.2. Let $\mathfrak{F}$ be an indexed set starting with $\varnothing$. $\mathfrak{F}$ is said to be a *logical cover* of the set Σ, relative to the logic X, indicated by $\text{COV}_X(\mathfrak{F}, \Sigma)$, provided:

[7] In saying this, we assume that nobody is obliged to bring about anything impossible and that whatever is self-inconsistent cannot truly be a belief.

[8] See especially [5, 6].

- For every element a of the indexed family, $\textsc{con}_X(a)$ and
- $\Sigma \subseteq \bigcup_{i\in I} \mathfrak{C}_X(a_i)$.

So an X-logical cover for Γ is an indexed family of sets starting with the empty set, such that there are enough logical resources in the cover to prove, in the logic X, each member of Γ. Evidently, given the rule [R], $\{\varnothing, \Gamma\}$ will always be a logical cover of Γ if the latter set is X-consistent, though it won't, in general, be the least.

If $\mathfrak{F}_\Sigma$ is a logical cover for the set Σ, the cardinality $|I| - 1$ where I is the index set for $\mathfrak{F}_\Sigma$, is referred to as the *width* of the cover, indicated by $w(\mathfrak{F}_\Sigma)$.

In the special circumstance that all the members of a logical cover of Σ are disjoint, the cover is said to *partition* Σ.[9]

And finally we introduce the notion at which we have hinted since the start of this section.

Definition 4.3. The *level* (relative to the logic X) of the set Γ of formulas of the underlying language of X, indicated by $\ell^X(\Gamma)$ is defined:

$$\ell^X(\Gamma) = \begin{cases} \min\limits_{w(\mathfrak{F})}[\textsc{cov}_X(\mathfrak{F}, \Gamma)] & \text{if this limit exists} \\ \infty & \text{otherwise} \end{cases}$$

In other words: the X-level (of incoherence or inconsistency) of a set Σ in a logic X is the width of the narrowest X-logical cover of Σ, if there is such a thing, and if there isn't, the level is set to the symbol ∞.

One might think that there will fail to be a narrowest logical cover when there is more than one – when several are tied with the least width, but this is a misreading of the definition. There might indeed be several distinct logical covers, but there can only be one least width (which they all share). The uniqueness referred to in the definition attaches to the width, not to the cover, so to speak.

The only circumstance in which there might fail to be a narrowest logical cover, is one in which Σ has no logical covers at all. In this circumstance Σ must contain what we earlier called an absurd formula.

This notion satisfies both the requirement that it distinguishes between inconsistent sets which contain absurd formulas, and those which don't, and the requirement that the predecessor notion of consistency relative to X appears as a special case. For it is clear that if Γ is an X-consistent set of formulas then a narrowest logical cover of Γ is $\{\varnothing, a_1\}$ where $\Gamma \subseteq \mathfrak{C}_X(a_1)$ and $\textsc{con}_X(a_1)$. So at least part of the earlier notion of $\textsc{con}_X(\Gamma)$ is captured by $\ell^X(\Gamma) = 1$.

There is even an interesting insight which comes out of this new idea. For there are *two* levels of X-consistency, 0 and 1. In our earlier naive approach, we thought of consistency as an entirely monolithic affair, but once we give the matter

[9]This should be contrasted with a covering family *being* a partition. We can recover the latter notion from this one by intersecting the covered set with each of the disjoint sets in the logical cover.

some thought we see that the empty set does indeed occupy a unique position in the panoply of X-consistent sets. If we know only that Γ and Σ are both X-consistent, nothing at all follows about the X-consistency of $\Gamma \cup \Sigma$. But we may rest assured that both $\Sigma \cup \varnothing$ and $\Gamma \cup \varnothing$ are X-consistent. And the same goes for the X-consequences of the empty set, namely X-theorems. By our definition, $\ell^X(\Delta) = 0$ if Δ is empty or any set of X-theorems, and of course such sets are consistent with any X-consistent set of formulas. It is tempting to call these level 0 sets, *hyperconsistent.*

5. Level preservation

So now that we have the concept of an X-level, should we be concerned about preserving such a thing? Perhaps there is no need for such concern, since it is at least possible that the logic X preserves its own level, isn't it? Well, in a word, no. It is not in general true that X preserves level beyond, of course the levels 0 and 1 of X-consistency. All the logics we consider not only do that, but are characterized by doing that.

Suppose the set Γ contains not only the contingent formula α but also a denial β of α although it does not contain any X-absurdities. Now, by definition of denial the pair set is inconsistent, and so $\mathbb{C}_X(\Gamma) = \mathbb{S}$. If there are X-absurdities then the X-level of the closure of Γ is ∞, i.e., $\ell^X(\mathbb{C}_X(\Gamma)) = \infty$.[10] This amounts to a massive failure to preserve level.

The obvious question to raise is this: given that ℓ^X is a generalization of $\textsc{con}_X$, is it the case that level characterizes logics in the same way that preserving consistency (in the strong sense) does? If not, then it seems that the generalization is not perhaps as central a notion as the root idea upon which it generalizes. Fortunately, for supporters of the general notion, we may prove the following generalization of the Generalized Consistency Theorem.

Theorem 5.1 (Level Characterization Theorem). *Suppose that X and Y are inference relations as in the Generalized Consistency Theorem over the same language $\mathcal{L}$, and let ℓ^X and ℓ^Y be the level functions associated with the respective inference relations. Then*

$$\left[\ell^X(\Gamma) = \ell^Y(\Gamma) \text{ for every set } \Gamma \text{ of formulas of } \mathcal{L}\right] \iff X = Y\,.$$

Proof. The proof depends upon the Generalized Consistency Theorem.

($\Longrightarrow$) Assume that $\ell_X(\Gamma) = \ell_Y(\Gamma)$ for every set Γ of formulas of $\mathcal{L}$. Then by definition the two agree on which sets have level 1 and level 0. But this is to say that X and Y agree on which sets are consistent. But by the Generalized Consistency Theorem, any two such logics (logics which we say are at evens) must be identical.

[10]The authors would like to thank the the anonymous referee for noticing this general argument.

($\Longleftarrow$) Assume that $X = Y$ and Suppose that for some arbitrary set Γ of formulas of $\mathcal{L}$ $\ell^X(\Gamma) = \xi$, ξ some cardinal. Then, by the definition, there is a narrowest X-logical cover $\mathfrak{F}_\Gamma$ such that $w(\mathfrak{F}_\Gamma) = \xi$. Since $X = Y$, it must be the case by the Generalized Consistency Theorem that the two logics agree on consistency (in the strong sense). Further, by definition each $a_i \in \mathfrak{F}_\Gamma$ is such that $\text{CON}_X(a_i)$. But then, since X and Y are at evens, $\text{CON}_Y(a_i)$. Thus, $\mathfrak{F}_\Gamma$ must be a Y-logical cover of Γ of width ξ. Moreover, this must be the narrowest such logical cover or else by parity of reasoning, there would be an X-logical cover of cardinality less than ξ contrary to hypothesis. Since Γ was arbitrary it follows that ℓ^X and ℓ^Y must agree on all sets of formulas of the language $\mathcal{L}$. □

This suggests that level is worth preserving, that it is a sort of natural logical kind, but doesn't show how the preservation may be carried out. It is time to repair that lack.

Perhaps the most straightforward route to preserving X-level is to define a new inference relation based on X. Evidently, the definition in question must also connect somehow with the notion of X-level and thus ultimately to CON_X (from now on we shall mostly drop reference to the background logic, like X, when no confusion will result). The process might have been informed by the ancient joke:

Question: How do you get down from an elephant?
Answer: You don't get down from an elephant, you get down from a duck.

except in our case the question and answer would go:

Question: How do you reason from inconsistent sets?
Answer: You don't reason from inconsistent sets, since every formula follows in that case, you reason from consistent *subsets*.

In other words, an inconsistent set is one for which the distinction between what follows and what doesn't has collapsed. This lack of meaningful contrast means that it no longer makes sense to talk about inferring conclusions from such a set. In order to regain the distinction we are going to have to drop back to the level of consistency and the only way to do that, is to look at consistent subsets of the original set.

Absent the notion of level, there are different ways to do this. The one suggested in [4] to deal with inconsistent sets of beliefs, involves two stages: At the first stage we discover the smallest subset of the inconsistent set which still exhibits the inconsistency. At the second stage we discard the member of the inconsistent subset with the least evidence, and repeat as necessary until the set is consistent. Having thus cleansed the belief set, we may now draw conclusions as we did before.[11]

We don't say that this process can't work. We do say that it doesn't seem to work in every case. There is a clear difficulty here when the two conditions

[11] This procedure was intended to apply to classical inference, but the method obviously generalizes to cover cases in which the base inference is X.

on rational belief: consistency (which we might call the external condition) and evidential support (the internal condition), pull us in different directions.

In the lottery paradox, for instance, we seem to have good evidence for each one of the lottery beliefs (ticket 1 won't win, ticket 2 won't win,..., ticket n won't win.) and we can make the evidence as strong as we like by making the lottery ever larger. Now conjoin the beliefs and we get 'No ticket will win.' which contradicts fairness. We could get consistency by throwing out the belief that the lottery is fair, but that would be cheating. The problem is that each of the lottery beliefs has exactly the same support as the others. They stand or fall as one, it would seem. If we let them all fall, then the rationality of buying a lottery ticket would seem to follow or at least the non-irrationality. But isn't it true that it isn't rational, according to the accepted canons at least, to buy a lottery ticket?[12]

Leaving aside the possibly controversial issue of the lottery paradox, take any situation in which we are unable to find a rationale for discarding one member of an inconsistent subset rather than another. Quine seems to suggest that in this situation, the counsel of prudence is to wait until we do find some way to distinguish among the problematic beliefs. Those with less patience seem to regard random discarding until at last we get to consistency, to be the path of wisdom.[13] We are inclined to reply to Quine that patience, for all that it is a virtue, is sometimes also a luxury we cannot afford or even a self-indulgence that we do well to deny ourselves.

To the others we say, consistency is not a virtue which trumps everything else. Suppose we might achieve consistency by throwing away one of α or β though we have no reason to prefer one over the other. Flipping a coin is a method for determining which goes to the wall, but we have no way of knowing if we have determined the correct one. We have left ourselves open to having rejected a truth or accepted a falsehood. 'Yes, but at least we now have consistency!' won't comfort us much if the consequences of picking the wrong thing to throw away are unpleasant enough.

Let us take up level once more.[14] In saying the level of the set Σ is k, we are saying two things. First that there is a way to divide the logical resources of Σ into k distinct subsets each of which is consistent. From now on we shall refer to these consistent subsets as *cells*. Second, that any way of thus dividing Σ must have at least k cells. Here we have got to the level of consistency, but we've got there k times. Not only might we wonder which of the k-cells is the 'real' one, the one which best represents 'the way things really are,' but there may be lots and lots of distinct ways to form the k cells. Which of the possibly many ways should we privilege?

[12]Not for nothing have lotteries long been known as 'a tax on fools.'

[13]This seems to be the route advocated some of those in computing science who have devised so-called truth-maintenance systems.

[14]We realize that the Quinean suggestion is not the only one, though it might be the most well-known, to deal with inconsistent sets of formulas. We do not, however, intend this essay as a survey of all of so-called paraconsistent logic.

At this point, we cannot answer these questions, which means that we must treat the cells on an equal footing along with the various ways of producing them.[15] In saying this, we say that we shall count as a consequence of Σ in the reconstructed inference relation, whatever formula follows (in the 'underlying' logic, say X) from at least one cell in every way of dividing Σ into k cells.

When the underlying logic is X, the derived inference relation is called X-level forcing, which relation is indicated by $[\Vdash_X$. We can give the precise definition as:

Definition 5.2. $\Gamma [\Vdash_X \alpha$ if and only if, for every division of Γ into $\ell^X(\Gamma)$ cells, for at least one of the cells $\Delta, \Delta \vdash_X \alpha$

It is easy to see that:

Proposition 5.3. *If* $\Gamma [\Vdash_X \alpha$ *then* $\ell^X(\Gamma) = \ell^X(\Gamma \cup \{\alpha\})$

Proof. Suppose the condition obtains and let $\ell^X(\Gamma) = k$. It follows from the definition that every division of Γ into k cells, results in at least one cell that X-proves α. But then we could add α to the cell in question without losing the cell property since X is a logic which preserves consistency. In such a case, after adding α we would have a division of $\Gamma \cup \{\alpha\}$ into k cells. Moreover there couldn't be a division of $\Gamma \cup \{\alpha\}$ into fewer than k cells without there being a similar division of Γ which would contradict the hypothesis. □

It obviously follows directly from this that:

Corollary 5.4. $[\Vdash_X$ *preserves X-level, in the sense that* $\ell^X(\Gamma) = \ell^X(\mathbb{C}_{[\Vdash_X}(\Gamma))$

6. Yes, but is it inference?

To be perfectly honest, or at least honest enough for practical purposes, all we have shown is that X-level forcing is a relation that preserves X-level. There is a gulf between this, and the assertion that $[\Vdash_X$ is an inference relation which preserves X-level. The obvious problem for anybody wishing to assert such a thing resides in the fact that we haven't, for all our efforts at precision, actually *said* which relations count as inference relations. What we *have* said, is that we assume that the inference relations we mention admit certain rules. Shall we take the collection of these rules to be constituitive of inference?

We shall not because some, at least, of the rules which the underlying logic admits, simply don't make sense for the derived relation. This should not come as a surprise. It is our palpable annoyance with the underlying logic which leads us to propose $[\Vdash_X$. How silly then to require that the derived logic inherit everything from the underlying logic since that would make the derived logic another source of irritation rather than the balm for which we hope.

[15] Which is not to say that there is no way. Elsewhere one might find suggestions which narrow the range of ways of dividing up our initial set. In this connection see the discussion of A-forcing in [6].

Although it is easy to check that $[\Vdash_X$ inherits from its underlying inference relation X both [R] and [Cut], we can see that it fails to admit the rule [Mon] of monotonicity, which we would do better to label the rule of *unrestricted* monotonicity from now on. But this is one of those cases in which the rule ought not to apply to the derived relation. If we are allowed to dilute premise sets in an arbitrary way, there is nothing to prevent us from raising the X-level of such sets. But raising the level gives us, in general, smaller cells in each logical cover. What used to X-follow from at least one cell of every such cover might no longer do so, as we are cut off from vital logical resources by the finer division.[16]

This is not to say that no form of monotonicity makes sense for the derived relation. Quite the contrary in fact, what most emphatically *does* make sense is that X-level forcing consequence must survive any dilution which preserves the level of the premise set. Such a restricted version of monotonicity manifestly is such a rule for X-level forcing as is trivial to verify. Along with level-preserving dilution, there are certain consequences which must survive any dilution at all, whether or not the X-level of the premise set increases. These are the consequences which dilution cannot affect, and we can say precisely which they are: the X-consequences of the empty set and of any unit set will remain X-consequences of at least one cell of every logical cover of any set which contains any of these privileged sets. In earlier work these sets were called *singular*.

The other properties which we have been mentioning for the underlying logics are non-triviality and having denial. It should be clear that when the underlying logic is non-trivial so will be its derived forcing relation. In fact, keeping to our original definition of consistency, in passing from an underlying logic X to its derived $[\Vdash_X$, many of the sets which are X-inconsistent fail to be $[\Vdash_X$-inconsistent, which is after all, the whole point of the derived relation.

Which brings us to denial. If the underlying logic has denial, nothing follows about the derived forcing relation, which is not necessarily a bad thing. This is because in the underlying logic, inconsistent sets are (typically) relatively easy to come by, but in the derived logic, the only inconsistent sets have inconsistent unit subsets, or what we called X-absurd formulas. Having denial doesn't imply having absurdities. So the derived logic will have denial only if the underlying logic has absurd formulas, but in no case will the derived logic have contingent denial.

And since β denies α in the derived logic if and only if one or both of the two are absurd in the underlying logic, the principle [Den] must hold of the derived relation but it is much less interesting there, than it is in the underlying logic.

So for the derived relation, we would seem to be on solid ground when we require [R], [Cut], and the restricted version of monotonicity. Those we might well regard as the hallmarks of inference, or at least of *derived* inference. And let us not forget that the derived relation agrees exactly with the underlying relation X

[16] Here is a concrete example where the underlying logic is classical. The premise set $\{\alpha, \alpha \supset \beta\}$ has the classical-level forcing consequence β since it has level 1. If we add the formula $\neg\beta$ the resulting set has level 2, and now there is a logical cover: $\{\varnothing, \{\alpha, \neg\beta\}, \{\alpha \supset \beta\}\}$ no cell of which classically proves β. Thus β is not a classical-level forcing consequence of the diluted set.

on the consequences of the X-consistent sets. So for this reason alone, we ought to admit the forcing relation into the fold.

Perhaps we should put it this way: Anybody who thinks that X is fine and dandy except for its failure to be properly sensitive to the varieties of inconsistent sets, *must* think that X-level forcing is an adequate account of inference. This is because, when premise sets are X-consistent X-level forcing just *is* X. And while it surely isn't X for (some) X-inconsistent sets, in those cases X isn't an inference relation. X has abdicated, throwing up its hands and retiring from the inferential struggle offering the hopeful reasoner nothing beyond a contemptuous 'Whatever!'

7. Forcing in comparison with other level-preserving relations

Finally, we consider the place of the X-level forcing relation compared with other possible relations which preserve X-level. We cannot claim uniqueness here, for there may be plenty of relations, even inference relations which preserve X-level. What we can claim however, is inclusiveness, in a sense to be made precise.

That precision will require another property[17] of the underlying logic X.

Definition 7.1. A logic X will be said to be *productival* if and only if for every finite set Γ there is some formula π such that

- $\pi \vdash_X \gamma$ for every $\gamma \in \Gamma$, and
- $\Gamma \vdash_X \pi$

Evidently being productival is another of those properties more honored at the level of underlying logics. If a productival logic X has denial, then X-level forcing will certainly not be productival. But of course at the underlying level, products are useful. For instance:

Theorem 7.2. *If X is productival then for any pair Γ, α with Γ a finite set of formulas and α a formula:*
If Y preserves X-level, admits level-preserving monotonicity [R] *and* [Cut]*, then* $\Gamma \vdash_Y \alpha \implies \Gamma[\Vdash_X \alpha$

Proof. We shall content ourselves with a sketch only – a fuller treatment can be found in [1]. Assume for indirect proof that $\Gamma \vdash_Y \alpha$ and that Γ fails to X-level force α. From the latter we know that Γ has finite level, say k, and that there is a logical cover of Γ of width k such that none of the k cells X-proves α. Where β (non-trivially) denies α add β to each cell and then form each of the k products of the cells. The set of these products when adjoined to the set Γ must have X-level k, but the Y closure of the set must have X-level k+1. So Y fails to preserve X-level contrary to hypothesis. □

[17] If we regard deductive systems as categories, then to call a logic productival is simply to say that the category (logic) X has products. The first condition amounts to the assertion of canonical projections while the second amounts to the universal mapping property of products.

The restriction to finite premise sets will chafe us only until we see its removal in the more general result referenced above.

So while there may be many inference relations which preserve X-level, X-level forcing is the largest of them.

References

[1] B. d' Entremont and G. Payette. Level compactness. *Notre Dame Journal of Formal Logic*, 47(4):545–555, December 2006.

[2] G. Gentzen. *The Collected Papers of Gerhard Gentzen.* Edited by M. E. Szabo. Studies in Logic and the Foundations of Mathematics. North-Holland Publishing Co., Amsterdam, 1969.

[3] A. Koslow. The implicational nature of logic: a structuralist account. In: *European Review of Philosophy*, Vol. 4, p. 111–155. CSLI Publ., Stanford, CA, 1999.

[4] W. V. O. Quine and J. S. Ullian. *The Web of Belief.* McGraw Hill, 2nd edition (1978), 1970.

[5] P. K. Schotch and R. E. Jennings. Inference and necessity. *Journal of Philosophical Logic*, 9:327–340, 1980.

[6] P. K. Schotch and R. E. Jennings On detonating. In: G. R. Priest, R. Routley, and J. Norman, editors, *Paraconsistent Logic*, p. 306–327. Philosophia Verlag, 1989.

Gillman Payette and Peter K. Schotch
Department of Philosophy
Dalhousie University
University Ave.
B3H 4P9 Halifax, Nova Scotia
Canada
e-mail: gpayette@dal.ca
peter.schotch@dal.ca

Submitted: 16 September 2006.
Accepted: 3 February 2007.