Truth and Permissive Consequence

Pablo Cobreros, Paul Égré, David Ripley, Robert van Rooij

[1]Universidad de Navarra

[2]Institut Jean Nicod

[3]University of Melbourne / University of Connecticut

[4]Institute of Logic, Language, and Computation

Amsterdam Workshop on Truth
March 14, 2013

Outline

Three challenges

Three well-known challenges for a Kripke-style theory of truth (viz. Field 2008, Beall 2011):

- Make the T-equivalences valid
- Obtain a well-behaved conditional
- Deal with the Strengthened Liar

Starting point: Truth and logical consequence

In bivalent classical logic, logical consequence is defined as the preservation of truth from premises to conclusions. Two equivalent definitions of $\Gamma \models \Delta$:

- the premises in Γ cannot be true without some conclusion in Δ being true
- the premises of Γ cannot be true without some conclusion in Δ being not false

Truth vs. Non-falsity

- General issue: what happens when moving to a larger truth value space, in which truth $\neq$ non-falsity?
- Answer: the two definitions can come apart
- Object of this paper: show that the second definition of logical consequence, from truth to non-falsity, gives rise to a nice and well-behaved logic in a trivalent setting.
- The resulting notion of consequence, which we may call permissive consequence, is fruitfully applied to the Liar and related paradoxes

Truth vs. Non-falsity

- General issue: what happens when moving to a larger truth value space, in which truth $\neq$ non-falsity?
- Answer: the two definitions can come apart
- Object of this paper: show that the second definition of logical consequence, from truth to non-falsity, gives rise to a nice and well-behaved logic in a trivalent setting.
- The resulting notion of consequence, which we may call permissive consequence, is fruitfully applied to the Liar and related paradoxes

Truth vs. Non-falsity

- General issue: what happens when moving to a larger truth value space, in which truth $\neq$ non-falsity?
- Answer: the two definitions can come apart
- Object of this paper: show that the second definition of logical consequence, from truth to non-falsity, gives rise to a nice and well-behaved logic in a trivalent setting.
- The resulting notion of consequence, which we may call permissive consequence, is fruitfully applied to the Liar and related paradoxes

Truth vs. Non-falsity

- General issue: what happens when moving to a larger truth value space, in which truth $\neq$ non-falsity?
- Answer: the two definitions can come apart
- Object of this paper: show that the second definition of logical consequence, from truth to non-falsity, gives rise to a nice and well-behaved logic in a trivalent setting.
- The resulting notion of consequence, which we may call permissive consequence, is fruitfully applied to the Liar and related paradoxes

Plan for this talk

- The Strict-Tolerant framework
- Transparent Truth
- Revenge issues

Outline

Permissive consequence: several precursors

- G. Malinowski (1990) on q-consequence and S. Frankowski (2004) on p-consequence: shifting the set of designated values between premises and conclusions, either by picking a sub- or a super-set.
- A. Nait-Abdallah 1995: studies the interplay between two notions of truth, classical truth (=1) and potential truth ($\neq 0$) in a trivalent setting
- B. Bennett 1998: introduces a notion of arguable entailment in a supervaluationist setting: "if all of the premises are 'unequivocally' true, then the conclusion is 'in some sense' true"
- E. Zardini 2008: main precursor of our work, defines consequence as going from designated values to tolerated values from premises to conclusion
- N. Smith 2008: fuzzy consequence defined as going from degrees > 0.5 to degrees ≥ 0.5 from premises to conclusions

Three-valued logic

In work we did on vagueness (Cobreros et al. 2012), we define strict and tolerant truth in terms of classical truth and a primitive notion of similarity. Here we define strict and tolerant truth directly within a three-valued framework, first by focusing on the language of first-order logic without identity:

- Truth-values: 0, 1/2, 1
- Connectives handled by means of Kleene's strong tables (conj: min, disj: max, neg: 1-).
- First-order models: $M = (D, I)$ where D domain of individuals, I function from predicates to $\{0, 1/2, 1\}$

Tolerant vs. Strict Interpretation

Framework: 3-valued logic, with Kleene/Priest valuation scheme:

Tolerant/Strict

- $M \models^s \phi$ iff $I(\phi) = 1$
- $M \models^t \phi$ iff $M(\phi) > 0$

Mixed consequence

The distinction between strict and tolerant values induces four notions of consequence:

mn-consequence

$\Gamma \models^{mn} \Delta$ iff if every formula γ in Γ is *m*-true, then some formula δ in Δ is *n*-true.

Overview of mixed consequence

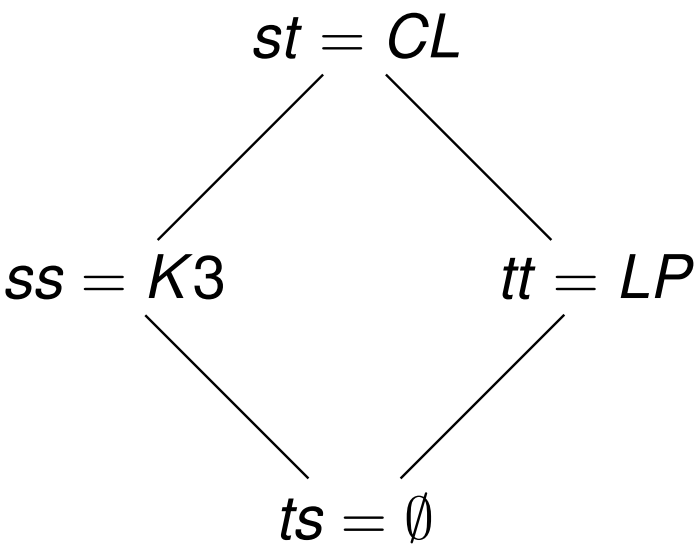

K3 vs. LP

The two 'unmixed' consequence relations (ss and tt) are well-known = Strong Kleene and LP. Note that each loses some desirable properties of the conditional and the deduction theorem:

- K3: $\phi \models_{K3} \phi$, but $\not\models_{K3} \phi \to \phi$
- LP: $\models_{LP} (\phi \wedge (\phi \to \psi)) \to \psi$ but $\phi \wedge (\phi \to \psi) \not\models_{LP} \psi$

ts vs. *st*

- *ts* can be viewed as restraining consequence: higher standard in conclusions than in the premises. It turns out to be so restraining as to be empty!
- *st* is permissive consequence: lower standard in conclusions than in the premises. The logic is just classical logic!

Remark about the conditional

$M \models^t A \to B$ iff: if $M \models^s A$ then $M \models^t B$

$M \models^s A \to B$ iff: if $M \models^t A$ then $M \models^s B$

		B		
	$\to$	1	$\frac{1}{2}$	0
	1	1	$\frac{1}{2}$	0
A	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
	0	1	1	1

Remark about the conditional

$M \models^t A \to B$ iff: if $M \models^s A$ then $M \models^t B$

$M \models^s A \to B$ iff: if $M \models^t A$ then $M \models^s B$

		B		
	$\to$	1	$\frac{1}{2}$	0
	1	1	$\frac{1}{2}$	0
A	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
	0	1	1	1

Remark about the conditional

$M \models^t A \to B$ iff: if $M \models^s A$ then $M \models^t B$

$M \models^s A \to B$ iff: if $M \models^t A$ then $M \models^s B$

		B		
	$\to$	1	$\frac{1}{2}$	0
	1	1	$\frac{1}{2}$	0
A	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
	0	1	1	1

Remark about the conditional

$M \models^t A \to B$ iff: if $M \models^s A$ then $M \models^t B$

$M \models^s A \to B$ iff: if $M \models^t A$ then $M \models^s B$

		B		
	$\to$	1	$\frac{1}{2}$	0
	1	1	$\frac{1}{2}$	0
A	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
	0	1	1	1

Remark about the conditional

$M \models^t A \to B$ iff: if $M \models^s A$ then $M \models^t B$
$M \models^s A \to B$ iff: if $M \models^t A$ then $M \models^s B$

		B		
	$\to$	1	$\frac{1}{2}$	0
	1	1	$\frac{1}{2}$	0
A	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
	0	1	1	1

Remark about the conditional

$M \models^t A \to B$ iff: if $M \models^s A$ then $M \models^t B$
$M \models^s A \to B$ iff: if $M \models^t A$ then $M \models^s B$

		B		
	$\to$	1	$\frac{1}{2}$	0
	1	1	$\frac{1}{2}$	0
A	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
	0	1	1	1

Remark about the conditional

$M \models^t A \to B$ iff: if $M \models^s A$ then $M \models^t B$
$M \models^s A \to B$ iff: if $M \models^t A$ then $M \models^s B$

		B		
	$\to$	1	$\frac{1}{2}$	0
	1	1	$\frac{1}{2}$	0
A	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
	0	1	1	1

Outline

Truth and Assertion

- So far we've talked of transparent vs. strict "truth", we actually prefer to distinguish two modes of assertion.
- The main idea is: regard logical consequence as reflecting constraints on assertion proper
- On the picture we are using: 1 and 0 do not primarily represent truth and falsity, but 'levels of goodness' (Zardini 2008), ways of marking that there is a nonarbitrary ground for asserting or denying a sentence. 1/2 will typically be used for sentences for which there is no nonarbitrary ground.

Incorporating truth

We now assume a first-order language with a distinguished truth-predicate T and a quote-forming operator $\langle\rangle$.

We use the same three-valued models, but with the following two constraints:

- $I(\langle A\rangle) = A$
- $I(T\langle A\rangle) = I(A)$ for every formula A. (identity for truth)

Transparent truth

We say that a truth predicate T is transparent iff A and $T\langle A\rangle$ are intersubstitutable everywhere without affecting validity.

Truth and consequences

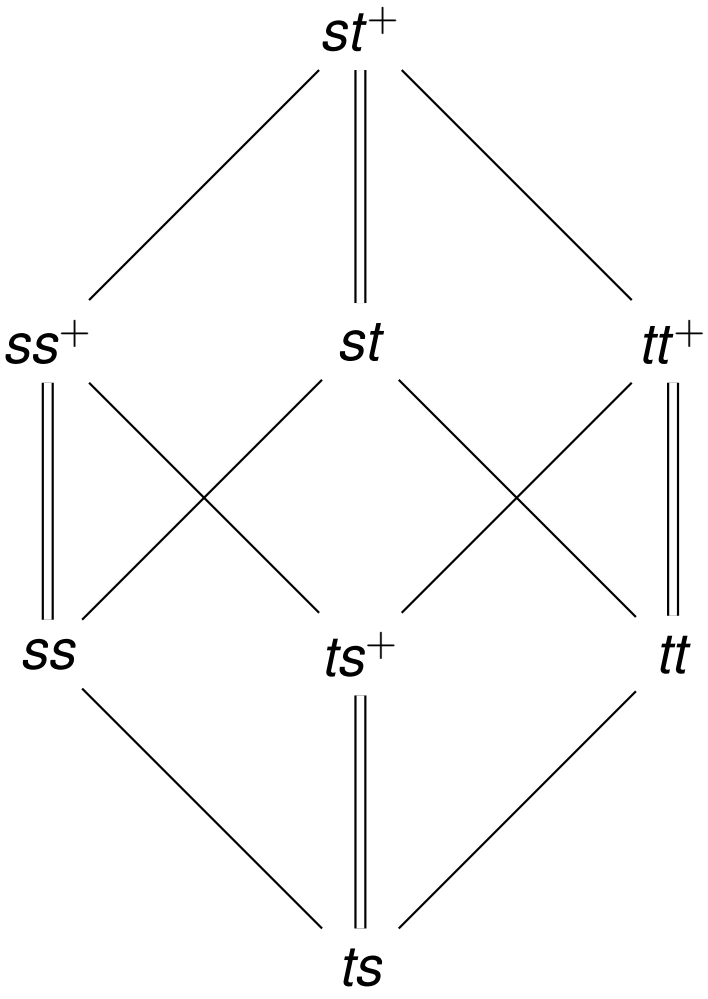

- Ordered by strength
- *ts* is almost empty, *ss* is K3, *tt* is LP, and *st* is classical logic.
- Double lines are conservative extensions (no change in T-free consequences).
- $^+$ consequence relations feature fully transparent T.

Truth and consequences

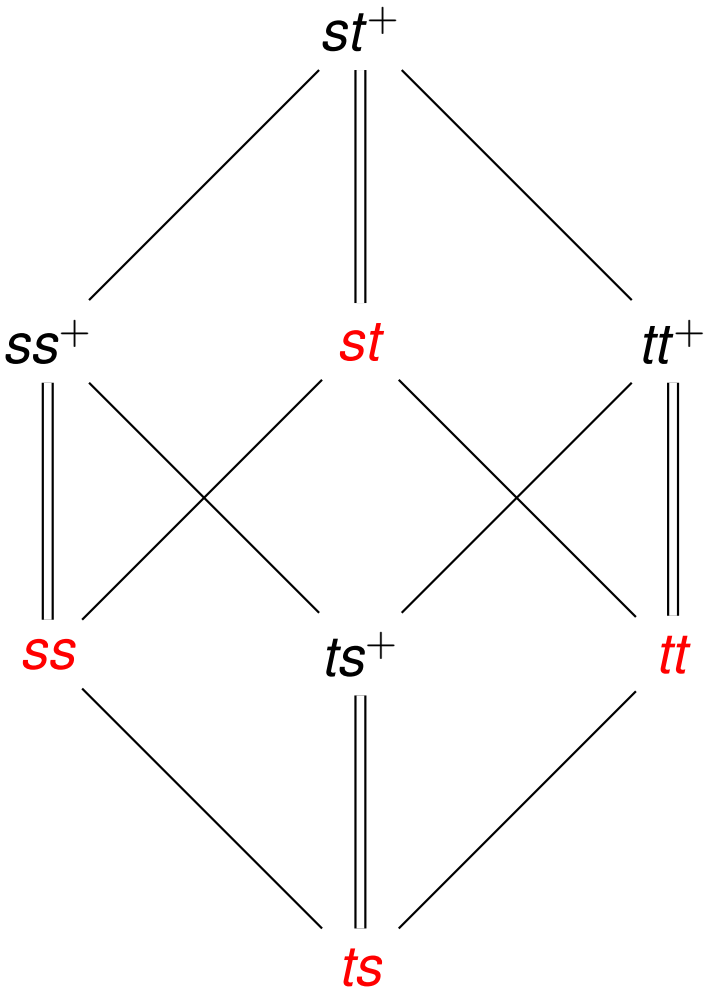

- Ordered by strength
- *ts* is almost empty, *ss* is K3, *tt* is LP, and *st* is classical logic.
- Double lines are conservative extensions (no change in T-free consequences).
- $^+$ consequence relations feature fully transparent T.

Truth and consequences

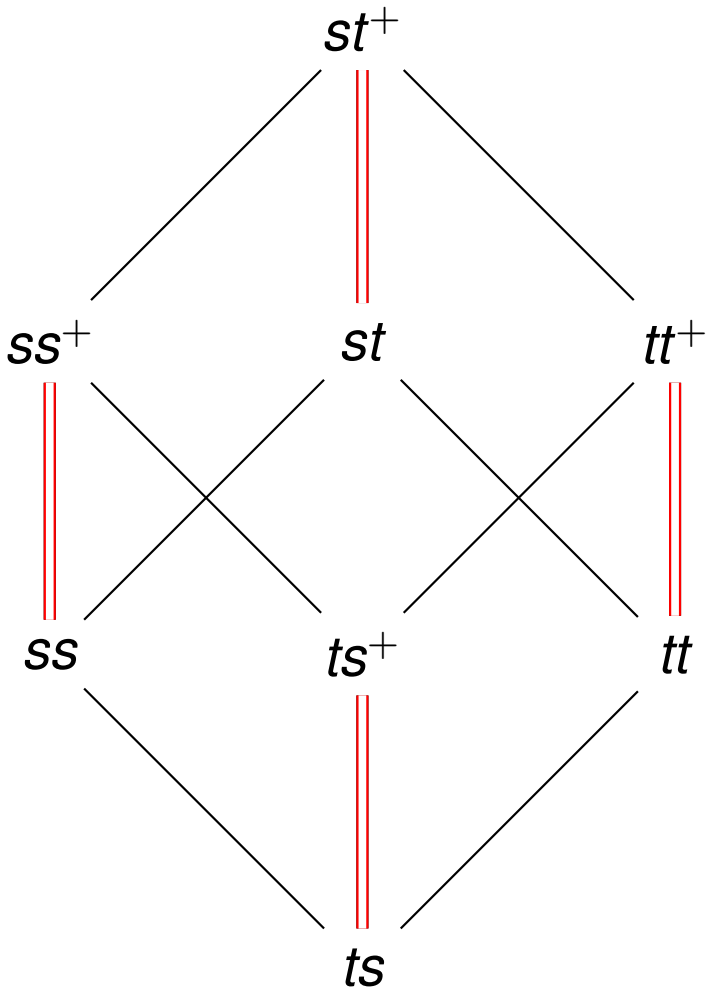

- Ordered by strength
- *ts* is almost empty, *ss* is K3, *tt* is LP, and *st* is classical logic.
- Double lines are conservative extensions (no change in T-free consequences).
- $^+$ consequence relations feature fully transparent T.

Truth and consequences

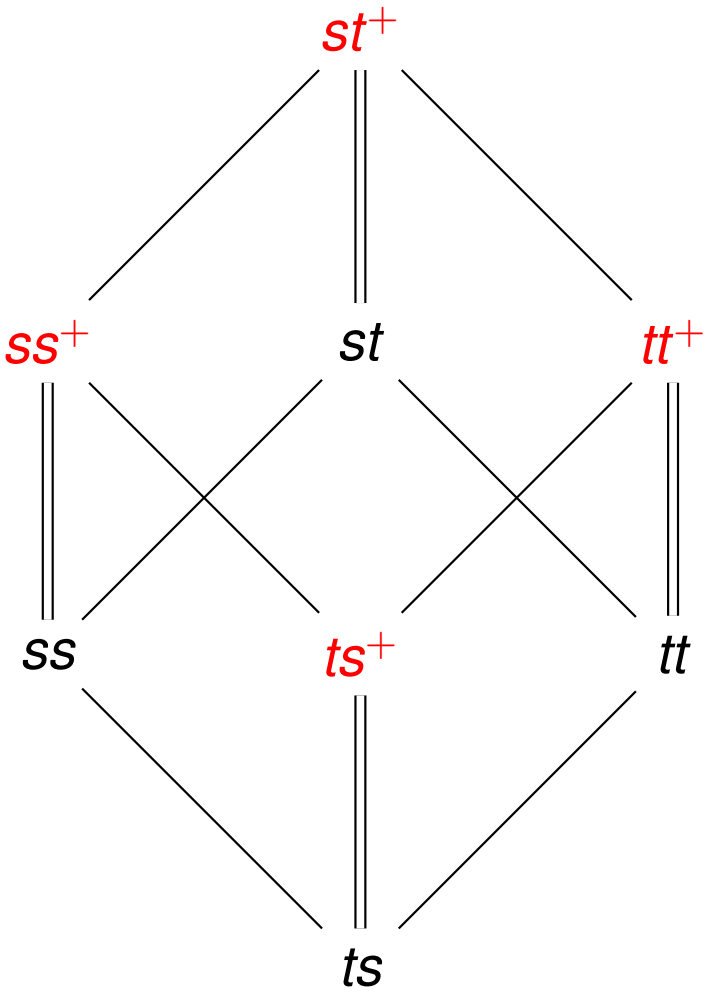

- Ordered by strength
- *ts* is almost empty, *ss* is K3, *tt* is LP, and *st* is classical logic.
- Double lines are conservative extensions (no change in T-free consequences).
- $^+$ consequence relations feature fully transparent T.

Main properties of the logic

We call STT the logic matching the $st+$ consequence relation:

- Conservative extension: STT is a conservative extension of classical logic
- Inference extension: if $\Gamma \vDash^{mn} \Delta$, then $\Gamma^\star \vDash_+^{mn} \Delta^\star$, where $\star$ is any uniform substitution.

Corollaries: T-schema and the conditional

STT validates all instances of the T-schema

STT incorporates a well-behaved conditional, satisfying the deduction theorem

Nontransitivity

The only price to pay for transparency is nontransitivity. Consider for instance the Liar:

Let $\lambda := \neg T\langle\lambda\rangle$. We have:

$$A \models^{STT} T\langle\lambda\rangle \wedge \neg T\langle\lambda\rangle$$

$$T\langle\lambda\rangle \wedge \neg T\langle\lambda\rangle \models^{STT} B$$

However, we do not have: $A \models^{STT} B$ for any A, B.

The Liar

$$
\dfrac{\dfrac{\dfrac{\top}{T\langle\lambda\rangle \vee \neg T\langle\lambda\rangle}\,\text{LEM} \qquad \dfrac{\dfrac{\dfrac{[T\langle\lambda\rangle]^1}{\lambda}\,\text{Transparency}}{\neg T\langle\lambda\rangle}\,\text{Def. }\lambda \qquad \dfrac{\dfrac{[\neg T\langle\lambda\rangle]^1}{\lambda}\,\text{Def. }\lambda}{T\langle\lambda\rangle}\,\text{Transparency}}{T\langle\lambda\rangle \wedge \neg T\langle\lambda\rangle}\,\wedge\text{I}}{T\langle\lambda\rangle \wedge \neg T\langle\lambda\rangle}\,\vee\text{E, 1}}{\bot}\,\text{Explosion}
$$

Outline

Four objections

Several objections can be made to our account:

- Objection 1: Is it not too high a price to give up transitivity?
- Objection 2: How can we know that a sentence is assertible only tolerantly and not strictly?
- Objection 3: How is the account going to help with revenge?
- Objection 4: isn't ST dealing with antinomies as fallacies of equivocation?

Objection1: The scope of nontransitivity

- All steps in the Liar reasoning are STT-valid, but chaining those steps is illegitimate
- Generalized transitivity: if $\Gamma \vDash^{STT} A, \Delta$ and $\Gamma, A \vDash^{STT} \Delta$ then $\Gamma \vDash^{STT} \Delta$
- We get a counterexample to generalized transitivity iff there is some model that assigns 1 to everything in Γ and 0 to everything in Δ, and every such model assigns $\frac{1}{2}$ to A.

Arguably therefore: nontransitivity is limited to pathological sentences

More about nontransitivity

- Nontransitivity plays the same role in our treatment of tolerance in the sorites paradox: every step is valid, but chaining them is illegitimate
- Both in our treatment of vagueness and of self-referential truth, transitivity can be retained everywhere where no special vocabulary is used, or where such vocabulary is used, but provided the standard of assertion can be kept constant

Objection 2: which sentences are strictly assertible?

Issue: why not abide by strict assertion only, and stick to transitive consequence?

We agree with Williamson and others that, as far as possible, assertion should go by strict standards

However, we want to be more liberal about assertion, by allowing the assertion or acceptance to go by the tolerant standard when there is no ground for strictly denying or rejecting a sentence.

Problem: are we always in a position to know that sentence is strictly assertible, or simply tolerantly assertible?

Answer: we agree with Kripke that, whether a sentence is Liar-like is contingent, and indeed, we cannot provide a context-independent characterization of which sentences are strictly assertible and which are not.

However, a strong Kleene adept is not in a better position than we are in this respect: hence the objection cannot be an argument for having a single mode of assertion, rather than two.

It would be an argument if a single-mode theory were performing equally on all other scores (see, however, S. Wintein's work and talk for more on this)

Objection 3: Revenge

Objection: how can the strict-tolerant account help with the Strengthened Liar?

Consider the sentence:

(1) this sentence is not strictly assertible

More formally

Let D be an operator such that: $M(DA) = 1$ if $M(A) = 1$, and $M(DA) = 0$ otherwise.

Let $\sigma := \neg DT\langle\sigma\rangle$

Then σ cannot receive any of the values 1, 0, 1/2 coherently.

A dilemma

There appears to be a basic dilemma for all theories of truth when it comes to revenge:

- Either limit the expressiveness of the language
- Or provide for indefinite extensibility of truth values [viz. Schlenker 2008, Cook 2008]

Comparing options

Option 1: declare σ illegitimate on grounds that assertibility is pragmatic, and not to be reflected in the language
Option 2: admit sentences like σ, and admit indefinite extensibility, but maintain that assertion and content should be distinguished

Option 2 appears to us to be more promising. One way of fleshing this out is to use a construction originally proposed by Priest (1984), and used by Ripley (2012) to deal with higher-order vagueness.

Hypercontradictions

Basic idea: start from $V_1 = \{1, 1/2, 0\}$, and consider $V_2 = \mathcal{P}(V_1) - \emptyset$, etc.

Declare a sentence to be tolerantly deniable in V_2 if its values all contain 0 or 1/2, tolerantly assertible if they all contain 1 or 1/2, strictly deniable if they contain only 0, and strictly assertible if they contain only 1.

If we collect the values σ would get over V_1, in V_2 we get: $\{0, 1\}$, or $\{1, 1/2, 0\}$

Hence σ is not strictly deniable, but it is tolerantly assertible: a good thing, since σ says it is not strictly deniable.

The construction can be generalized to further levels.

For each new level of truth values, we can construct a special predicate that will prevent the sentence from getting just one truth value in that set

But for each such predicate, we can collect which values it would get into a new set, and evaluate, relative to that higher-level, whether the sentence is tolerantly assertible or strictly assertible

Upshot: we get a hierarchy of truth values, but arguably, strict assertion and tolerant assertion hold through the entire hierarchy

Objection 4: equivocation

Objection: are we treating semantic antinomies as fallacies of equivocation?

In saying that the Liar is neither strictly assertible nor deniable, but tolerantly both, we may appear to declare the sentence ambiguous

But this is not what we say: we are not saying the Liar is ambiguous, only that whether it is assertible will depend on the mode of assertion.

Logic for assertoric ambiguity

> *"Logic for ambiguity? Who needs it? – I reply: pessimists". (D. Lewis 1982)*

Note that Lewis was considering logics for one-mode assertion. We keep the paraconsistent inspiration, but the logic is explicitly designed to deal with pragmatic or assertoric ambiguity.

Conclusions

Main benefits of the theory are:

- we retain Kripke's identity theory of truth
- we change the logic (make it more permissive) to satisfy transparency (hence the T-biconditionals)
- the change of logic gives us a well-behaved and simple conditional
- revenge issues are in a sense orthogonal to the choice of our logic: however we are in a position to acknowledge the problem, and can conciliate indefinite extensibility with the assertoric division between strict and tolerant
- Main difference with most theories of truth concerns the loss of transitivity: we hope to convince you that it is a feature, not a bug!

Related work

- Work by D. Ripley proposing a sequent calculus for STT
- Work by S. Wintein on tableaux for the strict-tolerant calculus incoporating a truth predicate
- Paper on which this talk is based: “Vagueness, truth and permissive consequence”, in which we bring side by side our treatment of truth and our treatment of vague sentences, with a unificatory purpose