Maximal consistent theories

Adolf Lindenbaum, 1904–1941.

Motivation: What can logic teach us about the world?

Logic is exclusively concerned with the form of an argument, not with its content.

Logic is exclusively concerned with the form of an argument, not with its content.

Suppose

$$p := \text{"The coin lands heads"} \; ,$$
$$q := \text{"The coin lands tails"} \; .$$

Logic is exclusively concerned with the form of an argument, not with its content.

Suppose

$$p := \text{"The coin lands heads"} \ ,$$
$$q := \text{"The coin lands tails"} \ .$$

Then it is intuitively obvious that "Either p or q" must be true.

Logic is exclusively concerned with the form of an argument, not with its content.

Suppose

$$p := \text{"The coin lands heads"}\ ,$$
$$q := \text{"The coin lands tails"}\ .$$

Then it is intuitively obvious that "Either p or q" must be true. So can we prove $\vdash p \vee q$?

Logic is exclusively concerned with the form of an argument, not with its content.

Suppose

$$p := \text{"The coin lands heads"}\ ,$$
$$q := \text{"The coin lands tails"}\ .$$

Then it is intuitively obvious that "Either p or q" must be true.
So can we prove $\vdash p \vee q$?
Of course not: $\nvdash p \vee q$ (by soundness). Our formal system does not know of our intended interpretation of p and q: they are just two atomic propositional formulæ.

Logic is exclusively concerned with the form of an argument, not with its content.

Suppose

$$p := \text{“The coin lands heads”}\ ,$$
$$q := \text{“The coin lands tails”}\ .$$

Then it is intuitively obvious that “Either p or q” must be true.
So can we prove $\vdash p \vee q$?
Of course not: $\nvdash p \vee q$ (by soundness). Our formal system does not know of our intended interpretation of p and q: they are just two atomic propositional formulæ.
Now set $S = \{p \leftrightarrow \neg q\}$. Now obviously $\vdash p \vee q$, under the additional set of assumptions S.

Logic is exclusively concerned with the form of an argument, not with its content.

Suppose

$$p := \text{“The coin lands heads”}\ ,$$
$$q := \text{“The coin lands tails”}\ .$$

Then it is intuitively obvious that “Either p or q” must be true.
So can we prove $\vdash p \vee q$?
Of course not: $\nvdash p \vee q$ (by soundness). Our formal system does not know of our intended interpretation of p and q: they are just two atomic propositional formulæ.
Now set $S = \{p \leftrightarrow \neg q\}$. Now obviously $\vdash p \vee q$, under the additional set of assumptions S.
Notation. $S \vdash \alpha \quad , \quad S \vdash_{\text{Ł}} \alpha$.

Question. What is it that such a set S of assumptions encodes? (Draw possible worlds at the board: only two possible worlds are left — $w(p) = 0$ and $w(q) = 1$, or conversely — out of the original 4 for p and q.)

Question. What is it that such a set S of assumptions encodes? (Draw possible worlds at the board: only two possible worlds are left — $w(p) = 0$ and $w(q) = 1$, or conversely — out of the original 4 for p and q.)

Answer. Your knowledge about the world. (In this case, that either the coin lands heads, or it lands tails, and *tertium non datur*.)

Question. What is it that such a set S of assumptions encodes? (Draw possible worlds at the board: only two possible worlds are left — $w(p) = 0$ and $w(q) = 1$, or conversely — out of the original 4 for p and q.)
Answer. Your knowledge about the world. (In this case, that either the coin lands heads, or it lands tails, and *tertium non datur*.)

Question. Does S encode complete knowledge about the world?[1]

Question. What is it that such a set S of assumptions encodes? (Draw possible worlds at the board: only two possible worlds are left — $w(p) = 0$ and $w(q) = 1$, or conversely — out of the original 4 for p and q.)
Answer. Your knowledge about the world. (In this case, that either the coin lands heads, or it lands tails, and *tertium non datur*.)

Question. Does S encode complete knowledge about the world?[1]
Answer. No, because the knowledge encoded by S can be properly increased without precipitating inconsistency.

Question. What is it that such a set S of assumptions encodes? (Draw possible worlds at the board: only two possible worlds are left — $w(p) = 0$ and $w(q) = 1$, or conversely — out of the original 4 for p and q.)
Answer. Your knowledge about the world. (In this case, that either the coin lands heads, or it lands tails, and *tertium non datur*.)

Question. Does S encode complete knowledge about the world?[1]
Answer. No, because the knowledge encoded by S can be properly increased without precipitating inconsistency. (For example, you may know that p is necessarily the case, perhaps because the coin is fake and bears two heads, or because we are talking about a toss that actually happened yesterday, or because you foresee the future.)

Question. What is it that such a set S of assumptions encodes? (Draw possible worlds at the board: only two possible worlds are left — $w(p) = 0$ and $w(q) = 1$, or conversely — out of the original 4 for p and q.)
Answer. Your knowledge about the world. (In this case, that either the coin lands heads, or it lands tails, and *tertium non datur*.)

Question. Does S encode complete knowledge about the world?[1]
Answer. No, because the knowledge encoded by S can be properly increased without precipitating inconsistency. (For example, you may know that p is necessarily the case, perhaps because the coin is fake and bears two heads, or because we are talking about a toss that actually happened yesterday, or because you foresee the future.)

Maximal consistent extension of S:

$$S \cup \{p\} \quad , \quad S \cup \{q\} .$$

Now a more systematic development of these ideas.

Now a more systematic development of these ideas.

A theory in CL (or any logic) is any set of formulæ that is closed under provability, i.e. is deductively closed.

Now a more systematic development of these ideas.

A theory in CL (or any logic) is any set of formulæ that is closed under provability, i.e. is deductively closed.

For any $S \subseteq \mathrm{F{\small ORM}}$, the smallest theory that extends S exists: it is the deductive closure $S^{\vdash}$ of S, defined by $\alpha \in S^{\vdash}$ if, and only if, $S \vdash \alpha$.

Now a more systematic development of these ideas.

A theory in CL (or any logic) is any set of formulæ that is closed under provability, i.e. is deductively closed.

For any $S \subseteq \textsc{Form}$, the smallest theory that extends S exists: it is the deductive closure $S^{\vdash}$ of S, defined by $\alpha \in S^{\vdash}$ if, and only if, $S \vdash \alpha$.

In particular, then, $\emptyset^{\vdash} = \textsc{Thm}$.

Now a more systematic development of these ideas.

A theory in CL (or any logic) is any set of formulæ that is closed under provability, i.e. is deductively closed.

For any $S \subseteq \textsc{Form}$, the smallest theory that extends S exists: it is the deductive closure $S^{\vdash}$ of S, defined by $\alpha \in S^{\vdash}$ if, and only if, $S \vdash \alpha$.

In particular, then, $\emptyset^{\vdash} = \textsc{Thm}$.

A theory Θ is consistent if $\Theta \neq \textsc{Form}$, and inconsistent otherwise.

Now a more systematic development of these ideas.

A theory in CL (or any logic) is any set of formulæ that is closed under provability, i.e. is deductively closed.

For any $S \subseteq \textsc{Form}$, the smallest theory that extends S exists: it is the deductive closure $S^{\vdash}$ of S, defined by $\alpha \in S^{\vdash}$ if, and only if, $S \vdash \alpha$.

In particular, then, $\emptyset^{\vdash} = \textsc{Thm}$.

A theory Θ is consistent if $\Theta \neq \textsc{Form}$, and inconsistent otherwise.

The theory Θ is maximally consistent, or maximal consistent, or even just maximal, if it is consistent, and whenever $\alpha \in \textsc{Form}$ is such that $\alpha \notin \Theta$, then $(\Theta \cup \{\alpha\})^{\vdash} = \textsc{Form}$.

It is customary to extend this terminology to arbitrary sets of formulæ.

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

Further, S is maximal consistent if it is consistent, and whenever $\alpha \in \textsc{Form}$ is not in $S^{\vdash}$, then $S \cup \{\alpha\}$ is inconsistent.

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

Further, S is maximal consistent if it is consistent, and whenever $\alpha \in \textsc{Form}$ is not in $S^{\vdash}$, then $S \cup \{\alpha\}$ is inconsistent.

Easy Facts.

(In classical logic.) For any $S \subseteq \textsc{Form}$, TFAE:

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

Further, S is maximal consistent if it is consistent, and whenever $\alpha \in \textsc{Form}$ is not in $S^{\vdash}$, then $S \cup \{\alpha\}$ is inconsistent.

Easy Facts.

(In classical logic.) For any $S \subseteq \textsc{Form}$, TFAE:

1. S is consistent.

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

Further, S is maximal consistent if it is consistent, and whenever $\alpha \in \textsc{Form}$ is not in $S^{\vdash}$, then $S \cup \{\alpha\}$ is inconsistent.

Easy Facts.

(In classical logic.) For any $S \subseteq \textsc{Form}$, TFAE:

1. S is consistent.
2. $S \nvdash \bot$.

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

Further, S is maximal consistent if it is consistent, and whenever $\alpha \in \textsc{Form}$ is not in $S^{\vdash}$, then $S \cup \{\alpha\}$ is inconsistent.

Easy Facts.

(In classical logic.) For any $S \subseteq \textsc{Form}$, TFAE:

1. S is consistent.
2. $S \nvdash \bot$.

Also, TFAE:

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

Further, S is maximal consistent if it is consistent, and whenever $\alpha \in \textsc{Form}$ is not in $S^{\vdash}$, then $S \cup \{\alpha\}$ is inconsistent.

Easy Facts.

(In classical logic.) For any $S \subseteq \textsc{Form}$, TFAE:

1. S is consistent.
2. $S \nvdash \bot$.

Also, TFAE:

1. S is maximal consistent.

It is customary to extend this terminology to arbitrary sets of formulæ.

Thus, $S \subseteq \textsc{Form}$ is consistent if $S^{\vdash} \neq \textsc{Form}$, and inconsistent otherwise.

Further, S is maximal consistent if it is consistent, and whenever $\alpha \in \textsc{Form}$ is not in $S^{\vdash}$, then $S \cup \{\alpha\}$ is inconsistent.

Easy Facts.

(In classical logic.) For any $S \subseteq \textsc{Form}$, TFAE:

1. S is consistent.
2. $S \nvdash \bot$.

Also, TFAE:

1. S is maximal consistent.
2. For any $\alpha \in \textsc{Form}$, either $S \vdash \alpha$ or $S \vdash \neg\alpha$, but not both.

The following item of historical interest was proved by Lindenbaum.

---

[2] *Added in Proof.* The original version of these slides contained an example which I said would illustrate the fact that non-constructive principles can entail ontological assumptions, such as "There exists at least one possible world". The example was wrong: its correction needs a subtle adjustment which I prefer not to discuss here, given that the whole discussion is an aside. I thank Tadeusz Litak for helping me clarify the issues involved, after my talk.

The following item of historical interest was proved by Lindenbaum.

Lindenbaum's Lemma.

(In classical logic.) For any consistent set $S \subset \mathrm{F{\scriptsize ORM}}$, there exists a maximal consistent theory Θ such that $\Theta \supseteq S$.

[2] *Added in Proof.* The original version of these slides contained an example which I said would illustrate the fact that non-constructive principles can entail ontological assumptions, such as "There exists at least one possible world". The example was wrong: its correction needs a subtle adjustment which I prefer not to discuss here, given that the whole discussion is an aside. I thank Tadeusz Litak for helping me clarify the issues involved, after my talk.

The following item of historical interest was proved by Lindenbaum.

Lindenbaum's Lemma.

(In classical logic.) For any consistent set $S \subset \textsc{Form}$, there exists a maximal consistent theory Θ such that $\Theta \supseteq S$.

This lemma is non-constructive; its proof uses the Axiom of Choice.[2]

---

[2] *Added in Proof.* The original version of these slides contained an example which I said would illustrate the fact that non-constructive principles can entail ontological assumptions, such as "There exists at least one possible world". The example was wrong: its correction needs a subtle adjustment which I prefer not to discuss here, given that the whole discussion is an aside. I thank Tadeusz Litak for helping me clarify the issues involved, after my talk.

There is a semantic counterpart to $S \vdash \alpha$ and $S^{\vdash}$.

There is a semantic counterpart to $S \vdash \alpha$ and $S^{\vdash}$.

If $S \subseteq \textsc{Form}$ is any set, a set α is a semantic consequence of S if any assignment $w\colon \textsc{Form} \to \{0, 1\}$ such that $w(S) = \{1\}$ is such that $w(\alpha) = 1$.

There is a semantic counterpart to $S \vdash \alpha$ and $S^{\vdash}$.

If $S \subseteq \textsc{Form}$ is any set, a set α is a semantic consequence of S if any assignment $w\colon \textsc{Form} \to \{0, 1\}$ such that $w(S) = \{1\}$ is such that $w(\alpha) = 1$.

We write $S^{\models}$ for the closure of S under semantic consequence.

Strong Completeness Theorem for CL

For any $\alpha \in \textsc{Form}$, and any set $S \subseteq \textsc{Form}$,

$$S \vDash \alpha \quad \text{if, and only if,} \quad S \vdash \alpha \,.$$

That is,

$$S^{\vDash} = S^{\vdash} \,.$$

(Similarly for n variables.)

Three key points about the rôle of logic.

Three key points about the rôle of logic.

1. Logic can teach us nothing (factual).

Three key points about the rôle of logic.

1. Logic can teach us nothing (factual).

2. Logic can model the factual (synthetic, extra-logical) knowledge that an agent already has about the world by encoding it into a consistent theory.

Three key points about the rôle of logic.

1. Logic can teach us nothing (factual).

2. Logic can model the factual (synthetic, extra-logical) knowledge that an agent already has about the world by encoding it into a consistent theory.

3. Maximal consistent theories then precisely encode complete knowledge of an agent about the world, and they determine the (unique) world wherein the agent is.

The notions related to theories that we have introduced are immediately adapted to Łukasiewicz logic. So we can talk about:

The notions related to theories that we have introduced are immediately adapted to Łukasiewicz logic. So we can talk about:

- Theories.

The notions related to theories that we have introduced are immediately adapted to Łukasiewicz logic. So we can talk about:

- Theories.
- Consistent and inconsistent theories.

The notions related to theories that we have introduced are immediately adapted to Łukasiewicz logic. So we can talk about:

- Theories.
- Consistent and inconsistent theories.
- Syntactic and syntactic consequences of arbitrary sets of formulæ.

The notions related to theories that we have introduced are immediately adapted to Łukasiewicz logic. So we can talk about:

- Theories.
- Consistent and inconsistent theories.
- Syntactic and syntactic consequences of arbitrary sets of formulæ.

Notation:

$$S \vDash_{Ł} \alpha \quad , \quad S^{\vDash_{Ł}} .$$

$$S \vdash_{Ł} \alpha \quad , \quad S^{\vdash_{Ł}} .$$

Łukasiewicz logic fails strong completeness.

Łukasiewicz logic fails strong completeness.

Let $S \subseteq \textsc{Form}_1$ be the set of formulæ:

Łukasiewicz logic fails strong completeness.

Let $S \subseteq \textsc{Form}_1$ be the set of formulæ:

$$\varphi_n(p) := ((n+1)(p^n \wedge \neg p)) \oplus p^{n+1}\,,$$

for each integer $n \geqslant 1$, where

Łukasiewicz logic fails strong completeness.

Let $S \subseteq \textsc{Form}_1$ be the set of formulæ:

$$\varphi_n(p) := ((n+1)(p^n \wedge \neg p)) \oplus p^{n+1} \,,$$

for each integer $n \geqslant 1$, where

$$p^k := \underbrace{p \odot \cdots \odot p}_{k \text{ times}} \,,$$

$$kp := \underbrace{p \oplus \cdots \oplus p}_{k \text{ times}} \,.$$

Łukasiewicz logic fails strong completeness.

Let $S \subseteq \textsc{Form}_1$ be the set of formulæ:

$$\varphi_n(p) := ((n+1)(p^n \wedge \neg p)) \oplus p^{n+1},$$

for each integer $n \geqslant 1$, where

$$p^k := \underbrace{p \odot \cdots \odot p}_{k \text{ times}},$$

$$kp := \underbrace{p \oplus \cdots \oplus p}_{k \text{ times}}.$$

Then $S \nvdash_{Ł} p$, but $S \vDash_{Ł} p$.

To prove this we need to think of $\varphi \in \textsc{Form}_1$ as a function $\overline{\varphi}\colon [0,1] \longrightarrow [0,1]$.

To prove this we need to think of $\varphi \in \mathrm{F{\scriptstyle ORM}}_1$ as a function $\overline{\varphi}\colon [0,1] \longrightarrow [0,1]$.

Given $x \in [0,1]$, consider the (unique, by compositionality) p.w. $w_x\colon \mathrm{F{\scriptstyle ORM}}_1 \longrightarrow [0,1]$ such that $w_x(p) = x$, and set

$$\overline{\varphi}(x) = w_x(\varphi)\,.$$

To prove this we need to think of $\varphi \in \textrm{FORM}_1$ as a function $\overline{\varphi}\colon [0,1] \to [0,1]$.

Given $x \in [0,1]$, consider the (unique, by compositionality) p.w. $w_x\colon \textrm{FORM}_1 \to [0,1]$ such that $w_x(p) = x$, and set

$$\overline{\varphi}(x) = w_x(\varphi)\,.$$

Notation	**Formal semantics**
$\neg\alpha$	$w(\neg\alpha) = 1 - w(\alpha)$
$\alpha \wedge \beta$	$w(\alpha \wedge \beta) = \min\{w(\alpha), w(\beta)\}$
$\alpha \oplus \beta$	$w(\alpha \oplus \beta) = \min\{1, w(\alpha) + w(\beta)\}$
$\alpha \odot \beta$	$w(\alpha \odot \beta) = \max\{0, w(\alpha) + w(\beta) - 1\}$

Table: Formal semantics of connectives in Łukasiewicz logic.

$S \not\vdash_{\text{Ł}} p$, but $S \vDash_{\text{Ł}} p$.

(Proof-by-drawing, at the board.)

$$\varphi_n(p) := ((n+1)(p^n \wedge \neg p)) \oplus p^{n+1}\,.$$

Notation	Formal semantics
$\neg\alpha$	$w(\neg\alpha) = 1 - w(\alpha)$
$\alpha \wedge \beta$	$w(\alpha \wedge \beta) = \min\{w(\alpha), w(\beta)\}$
$\alpha \oplus \beta$	$w(\alpha \oplus \beta) = \min\{1, w(\alpha) + w(\beta)\}$
$\alpha \odot \beta$	$w(\alpha \odot \beta) = \max\{0, w(\alpha) + w(\beta) - 1\}$

Table: Formal semantics of connectives in Łukasiewicz logic.

$S \nvdash_{Ł} p$, but $S \vDash_{Ł} p$.

(Proof-by-drawing, at the board.)

$$\varphi_n(p) := ((n+1)(p^n \wedge \neg p)) \oplus p^{n+1}.$$

Taking stock. $\vdash_{Ł}$ is compact, but $\vDash_{Ł}$ is not.

Notation	Formal semantics
$\neg\alpha$	$w(\neg\alpha) = 1 - w(\alpha)$
$\alpha \wedge \beta$	$w(\alpha \wedge \beta) = \min\{w(\alpha), w(\beta)\}$
$\alpha \oplus \beta$	$w(\alpha \oplus \beta) = \min\{1, w(\alpha) + w(\beta)\}$
$\alpha \odot \beta$	$w(\alpha \odot \beta) = \max\{0, w(\alpha) + w(\beta) - 1\}$

Table: Formal semantics of connectives in Łukasiewicz logic.

$S \nvdash_{Ł} p$, but $S \models_{Ł} p$.
(Proof-by-drawing, at the board.)

$$\varphi_n(p) := ((n+1)(p^n \wedge \neg p)) \oplus p^{n+1} \,.$$

Taking stock. $\vdash_{Ł}$ is compact, but $\models_{Ł}$ is not.

Note. $S \vdash_{Ł} \alpha \Rightarrow S \models_{Ł} \alpha$ always.

Notation	Formal semantics
$\neg\alpha$	$w(\neg\alpha) = 1 - w(\alpha)$
$\alpha \wedge \beta$	$w(\alpha \wedge \beta) = \min\{w(\alpha), w(\beta)\}$
$\alpha \oplus \beta$	$w(\alpha \oplus \beta) = \min\{1, w(\alpha) + w(\beta)\}$
$\alpha \odot \beta$	$w(\alpha \odot \beta) = \max\{0, w(\alpha) + w(\beta) - 1\}$

Table: Formal semantics of connectives in Łukasiewicz logic.

Not all hope is lost.

Not all hope is lost.

Theories are sets of formulæ, though not arbitrary ones; it turns out that they can have arbitrarily bad recursion-theoretic properties, in general.

Not all hope is lost.

Theories are sets of formulæ, though not arbitrary ones; it turns out that they can have arbitrarily bad recursion-theoretic properties, in general.

E.g. a theory can be decidable, recursively enumerable but not decidable (=semidecidable), undecidable, etc.

Not all hope is lost.

Theories are sets of formulæ, though not arbitrary ones; it turns out that they can have arbitrarily bad recursion-theoretic properties, in general.

E.g. a theory can be decidable, recursively enumerable but not decidable (=semidecidable), undecidable, etc.

Those theories that can be "described by a finite amount of information" are especially important.

Not all hope is lost.

Theories are sets of formulæ, though not arbitrary ones; it turns out that they can have arbitrarily bad recursion-theoretic properties, in general.

E.g. a theory can be decidable, recursively enumerable but not decidable (=semidecidable), undecidable, etc.

Those theories that can be “described by a finite amount of information” are especially important.

A theory Θ is axiomatised by a set $S \subseteq \textsc{Form}$ of formulæ if it so happens that $\Theta = S^{\vdash}$; and Θ is finitely axiomatisable if S can be chosen finite.

Not all hope is lost.

Theories are sets of formulæ, though not arbitrary ones; it turns out that they can have arbitrarily bad recursion-theoretic properties, in general.

E.g. a theory can be decidable, recursively enumerable but not decidable (=semidecidable), undecidable, etc.

Those theories that can be "described by a finite amount of information" are especially important.

A theory Θ is axiomatised by a set $S \subseteq \textsc{Form}$ of formulæ if it so happens that $\Theta = S^{\vdash}$; and Θ is finitely axiomatisable if S can be chosen finite.

Exactly the same definitions apply to Łukasiewicz logic.

Completeness Theorem for f.a. theories in Ł

For any $\alpha \in \textsc{Form}$, and any finite set $F \subseteq \textsc{Form}$,

$$F \vDash_{Ł} \alpha \quad \text{if, and only if,} \quad F \vdash_{Ł} \alpha\,.$$

That is,

$$F^{\vDash_{Ł}} = F^{\vdash_{Ł}}\,.$$

(Similarly for n variables.)

Completeness Theorem for maximal theories in Ł

For any $\alpha \in \textsc{Form}$, and any maximal consistent set $M \subseteq \textsc{Form}$,

$$M \vDash_{\text{Ł}} \alpha \quad \text{if, and only if,} \quad M \vdash_{\text{Ł}} \alpha\,.$$

That is,

$$M^{\vDash_{\text{Ł}}} = M^{\vdash_{\text{Ł}}}\,.$$

(Similarly for n variables.)

Taking stock:
Semantic intuition about maximal consistent theories

Taking stock:
Semantic intuition about maximal consistent theories

If Θ is a m.c. theory in CL, then $\Theta \vdash X_i$ or $\Theta \vdash \neg X_i$ (but not both), so Θ uniquely determines a possible world.

Taking stock:
Semantic intuition about maximal consistent theories

If Θ is a m.c. theory in CL, then $\Theta \vdash X_i$ or $\Theta \vdash \neg X_i$ (but not both), so Θ uniquely determines a possible world.
Conversely: If $\Theta_w = \{\alpha \in \textsc{Form} \mid w(\alpha) = 1\}$, then Θ is a m.c. theory in CL.

Taking stock: Semantic intuition about maximal consistent theories

If Θ is a m.c. theory in CL, then $\Theta \vdash X_i$ or $\Theta \vdash \neg X_i$ (but not both), so Θ uniquely determines a possible world.
Conversely: If $\Theta_w = \{\alpha \in \textsc{Form} \mid w(\alpha) = 1\}$, then Θ is a m.c. theory in CL.

In classical logic, maximal consistent theories are the syntactic counterpart to possible worlds.

Artificial precision, revisited

Back to the philosopher's coat.

Artificial precision, revisited

Back to the philosopher's coat. Consider the vague proposition

$$X := X_1 : \text{“Phil's coat is red.”}$$

Artificial precision, revisited

Back to the philosopher's coat. Consider the vague proposition

$$\boxed{X := X_1 : \text{"Phil's coat is red."}}$$

Let us regard X as an atomic proposition in Łukasiewicz logic. Its intended semantics, or intended model, is the sentence in quotes – it is not a number, as you can plainly see.

Artificial precision, revisited

Back to the philosopher's coat. Consider the vague proposition

$$\boxed{X := X_1 : \text{“Phil's coat is red.”}}$$

Let us regard X as an atomic proposition in Łukasiewicz logic. Its intended semantics, or intended model, is the sentence in quotes – it is not a number, as you can plainly see.

To focus on the core of the matter, let us restrict attention to Łukasiewicz logic over the one variable X.

Artificial precision, revisited

Back to the philosopher's coat. Consider the vague proposition

$$\boxed{X := X_1 \,:\, \text{“Phil's coat is red.”}}$$

Let us regard X as an atomic proposition in Łukasiewicz logic. Its intended semantics, or intended model, is the sentence in quotes – it is not a number, as you can plainly see.

To focus on the core of the matter, let us restrict attention to Łukasiewicz logic over the one variable X.

Just as in the previous case of tossing a coin, if all we know about X is that it is a propositional variable, the story is over.

Artificial precision, revisited

Back to the philosopher's coat. Consider the vague proposition

$$\boxed{X := X_1 \ : \ \text{“Phil's coat is red.”}}$$

Let us regard X as an atomic proposition in Łukasiewicz logic. Its intended semantics, or intended model, is the sentence in quotes – it is not a number, as you can plainly see.

To focus on the core of the matter, let us restrict attention to Łukasiewicz logic over the one variable X.

Just as in the previous case of tossing a coin, if all we know about X is that it is a propositional variable, the story is over.

For the only formulæ $\alpha(X)$ that will be provable are analytic truths (relative to Łukasiewicz logic), which by their very nature are absolutely uninformative about the colour of Phil's coat.

At the other extreme, let us assume that we have complete knowledge of the contingent facts $\alpha(X)$ concerning X that hold in the intended model.

At the other extreme, let us assume that we have complete knowledge of the contingent facts $\alpha(X)$ concerning X that hold in the intended model.

That means that we are given a maximally consistent theory Θ over Łukasiewicz logic.

At the other extreme, let us assume that we have complete knowledge of the contingent facts $\alpha(X)$ concerning X that hold in the intended model.

That means that we are given a maximally consistent theory Θ over Łukasiewicz logic.

If Łukasiewicz logic indeed is a logic of vagueness, then the maximally consistent theory Θ is to be thought of as a complete precisification of our intended (vague) model, namely, of the English sentence “Phil’s coat is red.”

At the other extreme, let us assume that we have complete knowledge of the contingent facts $\alpha(X)$ concerning X that hold in the intended model.

That means that we are given a maximally consistent theory Θ over Łukasiewicz logic.

If Łukasiewicz logic indeed is a logic of vagueness, then the maximally consistent theory Θ is to be thought of as a complete precisification of our intended (vague) model, namely, of the English sentence "Phil's coat is red."

But where could such a maximal consistent theory Θ come from?

It comes from the extra-logical assumption

"'Phil's coat coat is red' is true to degree $r \in [0, 1]$"

It comes from the extra-logical assumption

"'Phil's coat coat is red' is true to degree $r \in [0, 1]$"

Specifically, this is a semantic assumption: it tells us that certain states of affairs, while perhaps logically consistent, are known (or assumed) not to be the case.

It comes from the extra-logical assumption

"'Phil's coat coat is red' is true to degree $r \in [0, 1]$"

Specifically, this is a semantic assumption: it tells us that certain states of affairs, while perhaps logically consistent, are known (or assumed) not to be the case.

It is reasonable to expect that the assumption is maximally strong, falling short only of the strongest, inconsistent assumption according to which everything is the case. For observe that the stronger an assumption is, the fewer models it has, i.e. the fewer are the possible worlds that are consistent with it. Now the assumption "'My coat is red' is true to degree r" leaves us with just one possible world consistent with it, namely, the one world in which my coat is red to degree exactly r.

All this is mathematically summarised as follows:

$$\Theta_r = \{\alpha(X) \in \textsc{Form}_1 \mid w_r(\alpha(X)) = 1\},$$

where $w_r \colon \textsc{Form}_1 \to [0,1]$ is the only possible world such that $w_r(X) = r$.

All this is mathematically summarised as follows:

$$\Theta_r = \{\alpha(X) \in \textsc{Form}_1 \mid w_r(\alpha(X)) = 1\},$$

where $w_r\colon \textsc{Form}_1 \to [0,1]$ is the only possible world such that $w_r(X) = r$.

The formulæ in Θ_r that are not analytic truths are precisely those synthetic, factual truths about the colour of Phil's coat that the semantic assumption $w(X_1) = r$ entails, and that Łukasiewicz logic is able to express syntactically.

All this is mathematically summarised as follows:

$$\Theta_r = \{\alpha(X) \in \textsc{Form}_1 \mid w_r(\alpha(X)) = 1\},$$

where $w_r \colon \textsc{Form}_1 \to [0,1]$ is the only possible world such that $w_r(X) = r$.

The formulæ in Θ_r that are not analytic truths are precisely those synthetic, factual truths about the colour of Phil's coat that the semantic assumption $w(X_1) = r$ entails, and that Łukasiewicz logic is able to express syntactically.

In other words, the theory Θ_r attempts to encode our semantic assumption about Phil's coat at the syntactic level, with the formal linguistic resources of Łukasiewicz logic.

All this is mathematically summarised as follows:

$$\Theta_r = \{\alpha(X) \in \textsc{Form}_1 \mid w_r(\alpha(X)) = 1\},$$

where $w_r \colon \textsc{Form}_1 \to [0,1]$ is the only possible world such that $w_r(X) = r$.

The formulæ in Θ_r that are not analytic truths are precisely those synthetic, factual truths about the colour of Phil's coat that the semantic assumption $w(X_1) = r$ entails, and that Łukasiewicz logic is able to express syntactically.

In other words, the theory Θ_r attempts to encode our semantic assumption about Phil's coat at the syntactic level, with the formal linguistic resources of Łukasiewicz logic.

Fact: Θ_r is a maximal consistent theory.

Key Question. Is the semantic assumption "'Phil's coat coat is red' is true to degree r" precisely equivalent to the set of syntactic assumption Θ_r?

Key Question. Is the semantic assumption "'Phil's coat coat is red' is true to degree r" precisely equivalent to the set of syntactic assumption Θ_r?

Theorem (Proof reducible to Hölder's Theorem, 1901)

The correspondence

$$r \longmapsto \Theta_r$$

yields a bijection between maximal consistent theories in Łukasiewicz logic over one variable, and real numbers $r \in [0, 1]$.

Key Question. Is the semantic assumption "'Phil's coat coat is red' is true to degree r" precisely equivalent to the set of syntactic assumption Θ_r?

Theorem (Proof reducible to Hölder's Theorem, 1901)

The correspondence

$$r \longmapsto \Theta_r$$

yields a bijection between maximal consistent theories in Łukasiewicz logic over one variable, and real numbers $r \in [0, 1]$.

The innocent-looking Łukasiewicz axioms characterise the real numbers.

- $\Theta(1) = \{ X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 1 *if and only if Phil's coat is red.*)

- $\Theta(1) = \{ X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 1 *if and only if Phil's coat is red.*)
- $\Theta(0) = \{ \neg X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 0 *if and only if Phil's coat is not red.*)

- $\Theta(1) = \{ X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 1 *if and only if Phil's coat is red.*)
- $\Theta(0) = \{ \neg X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 0 *if and only if Phil's coat is not red.*)
- $\Theta(\frac{1}{2}) = \{ 2(X \wedge \neg X) \}^{\vdash}$.

- $\Theta(1) = \{\, X \,\}^{\vdash}$. (*"Phil's coat is red" is true to degree* 1 *if and only if Phil's coat is red.*)
- $\Theta(0) = \{\, \neg X \,\}^{\vdash}$. (*"Phil's coat is red" is true to degree* 0 *if and only if Phil's coat is not red.*)
- $\Theta(\frac{1}{2}) = \{\; 2(X \wedge \neg X) \;\}^{\vdash}$.
- $\Theta(\frac{2}{3}) = \{\; 3(\; (X \wedge \neg X) \wedge \neg(\, X \rightarrow (X \wedge \neg X)\,)\;)\; \}^{\vdash}$.

- $\Theta(1) = \{ X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 1 *if and only if Phil's coat is red.*)
- $\Theta(0) = \{ \neg X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 0 *if and only if Phil's coat is not red.*)
- $\Theta(\frac{1}{2}) = \{ \ 2(X \wedge \neg X) \ \}^{\vdash}$.
- $\Theta(\frac{2}{3}) = \{ \ 3(\ (X \wedge \neg X) \wedge \neg(X \rightarrow (X \wedge \neg X)) \ ) \ \}^{\vdash}$.
- $\Theta(r)$ is finitely axiomatisable if and only if r is a rational number.

- $\Theta(1) = \{ X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 1 *if and only if Phil's coat is red.*)
- $\Theta(0) = \{ \neg X \}^{\vdash}$. (*"Phil's coat is red" is true to degree* 0 *if and only if Phil's coat is not red.*)
- $\Theta(\frac{1}{2}) = \{ 2(X \wedge \neg X) \}^{\vdash}$. $(\ldots)$
- $\Theta(\frac{2}{3}) = \{ 3((X \wedge \neg X) \wedge \neg(X \rightarrow (X \wedge \neg X))) \}^{\vdash}$. $(\ldots)$
- $\Theta(r)$ is finitely axiomatisable if and only if r is a rational number.

Solving the problem of artificial precision completely means filling in the ellipses in natural language.

Thank you for your attention.