

Session 6: Beyond the Sentence

Modality, Intensionality, and Dynamics

James Pustejovsky

Brandeis University
Department of Computer Science

CS 135: Computational Semantics

September 15, 2026

Where We Are

Sessions 1–5 built a semantics that works. It is:

- **compositional** — meanings built from parts by typed application;
- **truth-conditional** — every sentence gets a value in a model;
- **computable** — over a finite domain, satisfaction is a program.

It is also, in two specific ways, **too small**.

Extensional	a predicate <i>is</i> its extension
--------------------	-------------------------------------

Sentence-level	each sentence is evaluated on its own
-----------------------	---------------------------------------

Today: what each restriction costs, and what the standard repairs look like. This is a survey — we will see the shape of two large literatures rather than work through either.

Today

- 1 The extensionality problem
- 2 Possible worlds; intensional types
- 3 Modality: *must*, *might*, *necessarily*
- 4 Propositional attitudes; *de re* and *de dicto*
- 5 The anaphora problem
- 6 Dynamic semantics: meaning as context change
- 7 Discourse Representation Theory
- 8 What carries into Part II

Nothing today is on Problem Set 1, and Sessions 12–24 do not use it. It is here because a semantics course that never mentions these boundaries would be misrepresenting the field.

A Predicate Is Its Extension

In our fragment, $\llbracket \text{dog} \rrbracket$ is the set of dogs. That is the whole story.

Now suppose, as a matter of biological fact, that exactly the same creatures have hearts and kidneys. Then

$$\llbracket \text{creature with a heart} \rrbracket = \llbracket \text{creature with a kidney} \rrbracket$$

Our semantics says these two predicates *mean the same thing*.

That is wrong, and demonstrably so. *Alice is looking for a creature with a heart* and *Alice is looking for a creature with a kidney* can differ in truth value — **so meaning must be finer-grained than extension**.

The example is Quine's. It is the standard demonstration that coextensive is not cosignificant.

Three Symptoms of the Same Disease

Phenomenon	What breaks
Coextensive predicates	identified, though they differ in meaning
Attitude reports	<i>Alice believes Bob left</i> gets no truth conditions
Modal claims	<i>Bob might have left</i> is about how things <i>could</i> be

All three want the same thing: **a way to talk about how things are not, but could be.**

Our model has one world. Everything true in it is settled, and everything false is simply false. There is no room for alternatives — and every phenomenon above requires alternatives.


The repair is to add them.

Possible Worlds

The move

Replace the single model with a **set of worlds** $\mathcal{W}$. Each $w \in \mathcal{W}$ is a complete way things might be — its own interpretation function I_w .

Denotation becomes **world-relative**: $\llbracket \alpha \rrbracket^w$.



... and so on

Extension of φ at w its truth value in that world

Intension of φ the *function* from worlds to truth values

A **proposition** is an intension of a sentence: equivalently, the *set of worlds where it is true*.

Intensional Types

Add a new basic type s for worlds. Then for any type τ , the type $\langle s, \tau \rangle$ is an **intension**.

Type	Denotes	Called
t	a truth value	extension of a sentence
$\langle s, t \rangle$	worlds $\rightarrow$ truth values	a proposition
$\langle e, t \rangle$	a set of individuals	extension of a predicate
$\langle s, \langle e, t \rangle \rangle$	worlds $\rightarrow$ sets	a property
e	an individual	
$\langle s, e \rangle$	worlds $\rightarrow$ individuals	an individual concept

Now *creature with a heart* and *creature with a kidney* can differ: same set in *our* world, different sets in some other — so their intensions differ, and meaning is intension.

The Queen of England is the standard individual concept: a different person depending on the world, and on the time.

The Cost

	Extensional	Intensional
Basic types	e, t	e, t, s
Type of a sentence	t	$\langle s, t \rangle$
Type of a noun	$\langle e, t \rangle$	$\langle s, \langle e, t \rangle \rangle$
Every denotation	as written	wrapped in a world argument

Every type in the fragment gains an s . Montague's PTQ does this uniformly, and the denotations become considerably harder to read.

This is why most working systems stay extensional. The intensional apparatus is well understood and expensive, and for question answering over a database there is only one world anyway. Knowing when you can afford to ignore it is part of the skill.

Frege's Puzzle

The oldest form of the problem, and still the cleanest.

The Morning Star is the Evening Star.

Both expressions refer to Venus. So extensionally, the sentence says

$\text{venus} = \text{venus}$

— a triviality, knowable *a priori*.

But the identity was an **astronomical discovery**. Someone can competently understand both terms and not know they corefer.

Frege (1892): expressions have both a *Bedeutung* (reference) and a *Sinn* (sense) — a “mode of presentation”. Possible-worlds intensions make sense precise: two terms differ in sense if they pick out different things in some world.

It is not the only way, and arguably not a fully satisfying one — see the next slide.

Where Possible Worlds Are Still Too Coarse

Intensions fix coextension. They do not fix everything.

Alice believes that $2 + 2 = 4$.

Alice believes that every group of prime order is cyclic.

Both complements are **necessarily true** — true at every possible world. So both denote the *same* proposition: the set of all worlds.

Our semantics therefore predicts that anyone who believes one believes the other. That is false: mathematical ignorance is possible.

The problem of logical omniscience. A possible-worlds account makes every agent believe all the consequences of what they believe. This is a genuine open problem, not a detail — and proposals to fix it (impossible worlds, structured propositions, hyperintensions) are all live.

Intensional Verbs

Some verbs create intensional contexts, and some do not. The difference is testable.

Verb	Example	Entails existence?
<i>find</i>	<i>Alice found a unicorn</i>	yes — extensional
<i>seek</i>	<i>Alice sought a unicorn</i>	no — intensional
<i>own</i>	<i>Alice owns a unicorn</i>	yes
<i>want</i>	<i>Alice wants a unicorn</i>	no

Two further diagnostics for intensionality:

- **Substitution failure.** If the morning star is the evening star, *Alice seeks the morning star* does not entail *Alice seeks the evening star*.
- **Scope ambiguity.** Intensional verbs give *de re/de dicto* readings; extensional ones do not.

The type of *seek* must take a *quantifier intension* rather than an entity: $\langle\langle s, \langle\langle e, t \rangle, t \rangle \rangle, \langle e, t \rangle \rangle$.

Montague's PTQ is largely an argument about getting this type right.

Modal Operators

With worlds in hand, modality is quantification over them.

Operator	Reading	Truth at w
$\Box\varphi$	necessarily φ	φ true at <i>all</i> worlds accessible from w
$\Diamond\varphi$	possibly φ	φ true at <i>some</i> world accessible from w

$\Box$ is $\forall$ over worlds; $\Diamond$ is $\exists$ over worlds. The duality is the same one from Session 3:

$$\neg\Box\varphi \equiv \Diamond\neg\varphi \qquad \neg\Diamond\varphi \equiv \Box\neg\varphi$$

Accessibility $R \subseteq \mathcal{W} \times \mathcal{W}$ says which worlds count as alternatives from where you stand. Different choices of R give different flavors of modality.

flavors of Modality

The same words express different modalities, distinguished by what R is.

flavor	Accessible worlds	Example
Epistemic	compatible with what is known	<i>Bob might be home</i>
Deontic	compatible with the rules	<i>You must file by Friday</i>
Alethic	metaphysically possible	<i>$2 + 2$ must equal 4</i>
Bouletic	compatible with someone's wishes	<i>I want it to rain</i>

Bob must be home is ambiguous between epistemic (I infer it) and deontic (he is required to be).

One semantics, parameterized. The modal operator is fixed; the accessibility relation varies with context. That is Kratzer's analysis, and it is the standard treatment.

Properties of R Become Logical Laws

Constraints on accessibility correspond exactly to modal axioms.

Property of R	Axiom	Reads
reflexive	$\Box\varphi \rightarrow \varphi$	what is necessary is true
transitive	$\Box\varphi \rightarrow \Box\Box\varphi$	necessity iterates
symmetric	$\varphi \rightarrow \Box\Diamond\varphi$	what is true is necessarily possible
serial	$\Box\varphi \rightarrow \Diamond\varphi$	necessity implies possibility

The empirical payoff. Epistemic modality is *not* reflexive in general — *Bob must be home* does not entail that he is. Deontic modality is not either: *you must file* does not entail that you will.

So a semantic framework can be tested against which entailments speakers accept. This is one of the cleaner places where formal semantics makes checkable predictions.

Propositional Attitudes

Alice believes that Bob left.

The problem. If *Bob left* denotes a truth value, then *believe* takes a truth value as argument — and every true complement would be interchangeable:

Alice believes $2 + 2 = 4$ $\equiv$ *Alice believes* *Bob left*

whenever both complements are true. Plainly wrong.

The repair. *Believe* takes a **proposition** — an intension — not a truth value:

$\llbracket \text{believe} \rrbracket : \langle \langle s, t \rangle, \langle e, t \rangle \rangle$

$\llbracket \text{Alice believes Bob left} \rrbracket^w = 1 \iff \text{Bob left is true at every world compatible with Alice's belief}$

An attitude verb is a modal operator with a subject-relative accessibility relation.

De Re and De Dicto

Alice is looking for a unicorn.

De dicto	there is a unicorn-seeking going on no particular unicorn; consistent with there being none
De re	there is a particular unicorn she seeks commits us to a unicorn existing

This is scope ambiguity again — but the quantifier now scopes over an *intensional* operator rather than another quantifier.

<i>de dicto</i>	$\text{seek}(a, \lambda P \exists x [\text{unicorn}(x) \wedge P(x)])$	quantifier inside
<i>de re</i>	$\exists x [\text{unicorn}(x) \wedge \text{seek}(a, x)]$	quantifier outside

Session 5's machinery — QR, storage, continuations — handles this too. The mechanisms were not built only for quantifier–quantifier interaction.

The Second Boundary

So far: one sentence at a time.

Pronouns Across Sentences

A student walked in. She sat down.

Interpret each sentence separately, as our fragment does:

A student walked in. $\rightsquigarrow \exists x[\text{student}(x) \wedge \text{walk-in}(x)]$

She sat down. $\rightsquigarrow \text{sit-down}(y)$ *y free!*

The pronoun cannot reach the quantifier. The $\exists x$ closed off at the end of the first sentence; y in the second is unbound, so the second sentence is not even a *sentence* in the Session 3 sense — it has no truth value.

And gluing them together does not help:

$\exists x[\text{student}(x) \wedge \text{walk-in}(x)] \wedge \text{sit-down}(y)$

y is still outside the scope of $\exists x$. Conjunction does not extend a quantifier's scope.

The Donkey Sentence

The same problem, inside a single sentence.

If a farmer owns a donkey, he beats it.

What it means: every farmer who owns a donkey beats every donkey he owns.

What compositional translation gives:

$$\exists x[\text{farmer}(x) \wedge \exists y[\text{donkey}(y) \wedge \text{own}(x, y)]] \rightarrow \text{beat}(x, y)$$

with x and y free in the consequent.

The puzzle. The indefinites are existential in isolation, but come out *universal* here. And the pronouns need to be bound from inside the antecedent, where the quantifier's scope does not reach. [Geach's example, 1962, and the problem that launched dynamic semantics.](#)

The Dynamic Turn

The idea

A sentence does not denote a truth value. It denotes a **function from contexts to contexts** — an instruction for updating what the discourse has established so far.

	Static	Dynamic
A sentence denotes	a truth value	a context change
Meaning is	truth conditions	context change potential
Composition is	function application	sequencing of updates

Truth does not disappear — a sentence is true if updating with it does not empty the context. But truth is now derived from something more basic.

Kamp (1981); Heim (1982); Groenendijk & Stokhof (1991).

Why This Solves the Problem

The key move: **an indefinite introduces a discourse referent that persists.**

Step	Update	Context after
—	—	$\{ \}$
<i>A student walked in.</i>	add x ; require student, walked-in	$\{x\}$
<i>She sat down.</i>	require x sat down	$\{x\}$

The referent x introduced by the first sentence is **still available** when the second arrives. The pronoun picks it up.

In static semantics, $\exists x$ has a scope that closes. In dynamic semantics, the existential *extends* its scope rightward across the discourse. That single change is what the framework buys.

Discourse Representation Theory

DRT gives this a notation: the **DRS box**.


x
student(x) walk-in(x) sit-down(x)

- The **top** compartment lists discourse referents — the variables in play.
- The **bottom** lists conditions on them.
- Processing a new sentence *adds* to the box rather than building a separate formula.

The pronoun in the third condition simply reuses the x already present. No scope problem, because there is no scope to escape.

Donkeys, Resolved

If a farmer owns a donkey, he beats it.



The truth condition for a conditional DRS: every way of satisfying the left box must extend to a way of satisfying the right.

That is where the universal force comes from. The indefinites are not universal quantifiers — they introduce referents, and the *conditional* quantifies universally over ways of assigning them.

The pronouns are bound because x and y in the left box are accessible from the right. Accessibility is structural, not scopal.

What Dynamic Semantics Costs

Gains	cross-sentential anaphora; donkey sentences; presupposition projection; a natural account of discourse structure
Costs	meanings are no longer truth conditions composition is no longer function application the elegant type system of Sessions 2–5 must be rebuilt

There is a reconciliation. Dynamic meanings can be given as functions in a suitably rich type system — and in the functional-programming presentation, context-change potential is naturally expressed with a *state monad*.

We are not pursuing that here; it is a substantial topic, and EU Chapter 8 does it properly if you want it. The point for today is that the dynamic turn is a real change of framework, not a patch.

Tense, in Passing

Worlds are not the only extra parameter a sentence needs.

Alice was a student.

This is not about how things are, but how they *were*. The standard treatment mirrors modality: add a set of **times** T with an ordering, and evaluate relative to a time as well as a world — $\llbracket \alpha \rrbracket^{w,t}$.

Operator	Reads	True at t iff
PAST φ	it was the case that φ	φ true at some $t' < t$
FUT φ	it will be the case that φ	φ true at some $t' > t$

Same shape as $\Box$ and $\Diamond$: quantification over an index set, with an accessibility relation — here, temporal precedence.

Aspect and its interaction with quantification are a large literature; TimeML is one computational descendant.

Presupposition

A third thing sentences do beyond asserting.

Alice stopped smoking.

Asserts	Alice does not smoke now
Presupposes	Alice smoked before

The diagnostic: presuppositions survive negation. *Alice did not stop smoking* still conveys that she smoked before. Assertions do not survive; presuppositions do.

Why this belongs in a dynamic framework. A presupposition is a *requirement on the input context* — the update is defined only if the context already contains the presupposed material. Natural when meanings are context-change operations; awkward otherwise.

Heim's dissertation (1982) treats presupposition and anaphora with the same machinery, which is the main argument that the dynamic turn earns its cost.

A Worked Update

A student walked in. She sat down. Every professor noticed her.

After processing	Referents	Conditions
—	$\emptyset$	—
<i>A student walked in.</i>	x	$\text{student}(x), \text{walk-in}(x)$
<i>She sat down.</i>	x	+ $\text{sit-down}(x)$
<i>Every professor noticed her.</i>	x	+ $\forall y[\text{prof}(y) \rightarrow \text{notice}(y, x)]$

Note that x persists across all three sentences — introduced once, referred to twice.

And note what did *not* happen. The universal quantifier in the third sentence introduces y , but y does **not** persist: **Every professor noticed her. He frowned.* is bad.

Indefinites introduce accessible referents; universals do not. That asymmetry is data, and a dynamic theory has to predict it — which it does, because the universal's referent is introduced inside a subordinate box.

Part I, Complete

Session	Added	Still missing
1	composition as application	types
2	types; what can combine	truth
3	models; satisfaction	a working fragment
4	the fragment, executable	quantifiers
5	generalized quantifiers; scope	intensions, discourse
6	worlds; context change	similarity

Each row closed the previous row's gap. The last one does not get closed inside Part I.

Add worlds, add discourse referents, add whatever you like — in every framework today, a word is still an arbitrary symbol. *dog* and *wolf* remain as unrelated as *dog* and *Tuesday*. That is what Part II is about.

What Carries Forward

Idea	Returns in	As
Intension as a function	—	not used again in this course
Modality as quantification	Session 24	epistemic states, alignment
Discourse referents	Session 11	role–filler binding
Variable reuse	Part IV	the copying problem

The last row is the one to remember. Donkey anaphora needs a variable used twice — once in the antecedent, once in the consequent.

In Part IV we will meet a framework where binding is the wedge product, and where $\mathbf{a} \wedge \mathbf{a} = 0$. **A structure that cannot bind the same thing twice cannot do donkey anaphora.** Whether that is repairable is open.

You have now seen the same limitation three times: as mu-abstraction in λ -DCS, as shadowing in the lambda calculus, and as donkey sentences here.

Summary

Concept	Key point
Extensionality	a predicate is its extension; coextensive becomes identical
Possible worlds	a set $\mathcal{W}$; denotation is world-relative
Intension	a function from worlds; <i>this</i> is meaning
Proposition	intension of a sentence; a set of worlds
Type s	added to the inventory; every type gains an s
$\Box$, $\Diamond$	$\forall$ and $\exists$ over accessible worlds
Accessibility	its properties are modal axioms; flavors of modality
Attitudes	take propositions, not truth values
<i>de re</i> / <i>de dicto</i>	scope over an intensional operator
Anaphora problem	$\exists$ closes; the pronoun cannot reach it
Dynamic semantics	meaning is context change potential
DRT	boxes; referents persist; accessibility is structural

Coming Up

Session 6A: λ -DCS — the logical tradition at Freebase scale, and what it still cannot do.

Session 6B: Semantic types as vectors — the bridge into Part II.

Session 7: The distributional hypothesis. Part II begins.

Reading: course slides. EU Chapter 8 if you want dynamic semantics properly; Portner's *Modality* is the standard short introduction to the first half.

Problem Set 1 is due Friday 25 September. Nothing from today is on it.

Questions?

Next: Session 6A — λ -DCS