Formal Semantics for Philosophers

Lecture Notes — 2016

Contents

Preface

These notes were written for my FORMAL SEMANTICS FOR PHILOSOPHERS class taught at the University of Edinburgh in the fall of 2012. They have been revised and expanded each year that I have taught the course again.

The notes are based in large part on other texts, in particular those listed below, and they are supposed to merely function as a supplement to these texts.

· Heim and Kratzer (1998)

· von Fintel and Heim (2007)

I am confident that the notes contain several errors and typos, so I recommend taking due care when using them.

Anders J. Schoubye
January, 2016

Chapter 1

Natural Language Analysis

1.1 Aspects of Natural Language

- Natural languages have several in principle separate components including the following:

    - **Syntax**
      The rules according to which lexical items enter into acceptable structural arrangements, i.e. whether various sentences are grammatical or not.
    - **Phonology and Phonetics**
      The relation between a structured sequence of lexical items and a sequence of sounds (phonology) and the physical production and auditory reception of those sound sequences (phonetics).
    - **Semantics**
      The meaning of sentences as a function of the meaning of the individual words and their order of combination. The meaning of individual words and morphemes, and their inferential relations. (*more below...*)
    - **Pragmatics**
      Meaning beyond *literal* meaning, e.g. various extra-linguistic processes used to determine what information is communicated (perhaps in addition to the literal content). (*more below...*)

1.1.1 Meaning and Truth Conditions

- Semantics is the branch of linguistics that studies **meaning**, in particular the meaning of sentences. This raises an immediate and obvious question, namely what is meaning?
- A fundamental and widely accepted assumption among formal semanticists is that the concept of (sentence) meaning should be explicated in terms of the notion of truth.

Meaning as Truth Conditions
The meaning of a sentence equals the conditions under which it is true, in short its *truth conditions*.

- To understand the meaning of the sentences in e.g. (1) and (2) is to understand what the world would have to be like for the sentences to be true.

(1) Grass is blue.
(2) Mitt Romney bought a pink donkey at the 2010 olympics.

- The truth conditional notion of meaning is quite narrow and thus fails to capture all kinds of meaning. For example, the meaning of some types of sentences cannot straightforwardly be explicated in terms of truth conditions.

(3) How old is Bob? (interrogative)
(4) Eat your carrots! (imperative)

- It would be nonsensical to ask whether (3) or (4) is true or false.
- Moreover, many things other than sentences can correctly be said to have *meaning*.
  - Suppose that you see smoke coming out of a dumpster. It seems natural to say that this (the smoke) *means* fire.
  - But, it is not as if the smoke is *true* or *false* — or that it makes sense to talk about the conditions under which the smoke is true or false.
  - Similarly, a 'no smoking' sign might be said to have a meaning, namely the instruction to refrain from smoking. However, it does not seem accurate to say that the *sign* has truth conditions. The sign (itself) is not true or false — signs are simply not the kinds of things that can be true or false.
- In short, many things can be said to have meaning even without having anything resembling truth conditions.
- We will therefore assume that truth conditions are properties of *declarative* sentences (and not *interrogative* or *imperative* sentences such as those above or signs etc.).

1.1.2 Semantics and Pragmatics

- The distinction between *semantics* and *pragmatics* is notoriously vague (and this distinction has been the subject of a prolonged debate over the past 30-40 years).
- When meaning is analyzed in terms of truth conditions, we often rely on intuitive judgments about what is and is not part of truth conditional meaning (i.e. what is part of the *literal* meaning).
- We do this to determine whether which parts of the relevant communicated information is part of the semantics (the literal meaning) or part of the pragmatics (the non-literal, extra-linguistic meaning).

- For example, an utterance might seem to convey information that is not intuitively part of what the speaker (literally) said. This kind of information belongs to *pragmatics*.
- Below are some examples of communicated information that is not intuitively part of the literal meaning (or what is sometimes called *what-is-said*).

1.1.3 Examples of Non-Truth Conditional Aspects of Meaning

▸ **Implicatures**

Implicature

- When a speaker asserts (5b) in response to (5a), the speaker (often) implies that he believes that the gas station around the corner is open. Following Grice (1989), this kind of implied meaning is referred to as an *implicature*.

(5) a. I'm out of petrol.
b. There's a gas station around the corner.

- There are several reasons to think that the implicature "the gas station is open" is not part of the truth conditional meaning (viz. the literal meaning). One reason is that that the implied content does not seem to have any impact on the truth conditions. That is, even if the gas station around the corner were closed, the speaker's utterance in (5b) would still seem true.
- Moreover, As Grice observed, content that is merely implicated can be explicitly "canceled" without this leading to some kind of intuitive inconsistency, cf. (6).

Implicature Cancellation

(6) There's a gas station around the corner, but I don't know if it is open.

- In contrast, notice that if a speaker attempts to "cancel" part of what is literally said, the result is intuitively contradictory.

(7) # There's a gas station around the corner, but I don't know if it's around the corner.

(8) # There's a gas station around the corner, but I don't know if it's a gas station.

- Here is another example of an implicature.

(9) Some of the students passed the exam.

- The sentence in (9) is naturally interpreted as conveying that some, but not all, of the students passed the exam. Again, this seems to be merely *implicated* rather than part of the truth conditional meaning. Again, notice that a speaker can felicitously cancel the implicature.

(10) Some of the students passed the exam. Perhaps they all passed?

Metaphor

▸ **Metaphor**

- Metaphors are (arguably) another case of sentences that convey information beyond the literal meanings. Consider for example the sentences below.

(11) It's raining men.
(12) Time is a thief.
(13) She's the apple of my eye.

- A speaker who asserts (11) is communicating that there is an abundance of male suitors. However, that is not what the sentence literally means.
- Similar observations could be made for e.g. (12) and (13).

Metonomy

▸ **Metonomy**

- Another example of non-literal meaning comes from cases of so-called metonymy.
- In cases of metonomy, an expression is used as an associative trigger of a different but closely related entity.
- For example, in the sentence in (14), the definite description 'The White House' is used as a kind of associative trigger for people who work at The White House. So, the sentence communicates that the people at The White House are concerned even though this is not its literal meaning.

(14) The White House is concerned about the recent developments.

Expressives / Intensifiers

▸ **Expressives**

- The words 'damn' and 'fuck' clearly convey something about the speaker's attitude towards the individual referred to, but they do not seem to contribute anything to the truth conditions of the sentences.

(15) That *damned* dog ate my homework.
(16) That *fucking* chef ruined the tuna.

- For example, as long as the dog did eat the speaker's homework, (16) is intuitively true.

1.1.4 Controversial Cases

There are many linguistic phenomena for which it is controversial as to whether they are semantic or pragmatic. One example is so-called *presuppositions*.

▸ **Presupposition**

- A sentence such as (17) implies (in some sense) that there is a president of England.
- Similarly, the sentence in (18) implies that Sue used to smoke.
- Notice that it would be inappropriate to assert either (17) or (18) if the speaker did not believe those implications.

(17) The president of England is bald.
(18) Sue stopped smoking.

- One strong reason for assuming that presuppositions are distinct from implicatures is that presuppositions are not (generally) cancellable. To see this, notice that the following sentences are infelicitous.

(19) # The president of England is bald, although I don't mean to imply that there is a president of England.
(20) # Sue stopped smoking, although I don't mean to imply that Sue ever smoked.

- This suggest that the "implications" generated by (17) and (18) are not just simple implicatures. Another option is to treat them as logical entailments?
- For example, (21) logically entails (22), because in every possible world where (21) is true, so is (22).

(21) Frank is bald
(22) Someone is bald.

- Perhaps these presuppositional implications are just like logical entailments, i.e. in every world where 'Sue stopped smoking' is true, so is 'Sue used to smoke'.
- The primary problem with this assumption is that presuppositions exhibit behavior quite distinct from logical entailments. For example, if (21) is embedded in the antecedent of a conditional, it no longer entails (22). The same happens if (21) is embedded under e.g. a possibility modal.

(23) If Frank is bald, he wears a hat.
$\nRightarrow$ Someone is bald.
(24) It's possible that Frank is bald.
$\nRightarrow$ Someone is bald.

- For this reason, *if*-clauses and possibility modals are often said to be *entailment-canceling* operators.
- In contrast, presuppositions intuitively survive these kinds of embeddings.

(25) If the president of England is bald, he probably wears a hat.
$\Rightarrow$ There is a president of England.
(26) It's possible that the president of England is bald.
$\Rightarrow$ There is a president of England.

- In short, there are good reasons to think that presupposition is a distinct linguistic phenomenon from both implicatures and entailments. The question is whether presuppositions are best analyzed as a semantic or a pragmatic phenomenon. We will discuss this in more detail later.

1.1.5 Truth Conditions and Compositionality

- The main focus here will be *compositional* semantics. Semantics is the branch of linguistics that deals with meaning, but what does 'compositional' mean?
- Compositionality is a a fundamental feature of natural languages. It is sometimes referred to as Frege's Principle, as Frege was the first to articulate it. The principle of compositionality can be stated as follows.

> **The Principle of Compositionality**
> The meaning of any given sentence S is a function of the meaning of the parts of S and the way in which these are combined.

- There are many reasons to think that the principle of compositionality holds for natural languages.
- First, natural languages contain a finite number of lexical items and a finite number of combinatory (syntactic) rules. Yet, in every known natural language it is possible to construct a potential infinity of grammatical sentences.
- Think, for example, of expressions such relative clauses ('who kissed Frank'), attitude verbs ('Frank believes that'), conjunction ('and'). These allow repeated (recursive) applications each generating a new sentence with a distinct meaning, e.g. 'Frank believes that Frank believes that Frank believes that ϕ'.
- Each of these sentences are in principle interpretable by competent speakers. In other words, competent speakers who have clearly limited resources have the capacity to understand a potential infinity of sentences.
- The best explanation for this capability is that meaning is compositional. In other words, speakers are capable from a finite number of lexical items (e.g. words) and from a finite number of combinatory rules of deriving the truth conditions of any sentence.
- For this reason, we take it as a fundamental desideratum of any semantic theory that it is a compositional theory.

Chapter 2

Summary: First Order Logic

2.1 First Order Logic (FOL)

- We start today with a recap of the syntax and semantics of **First Order Logic** (FOL). However, first we need to provide a vocabulary (or a lexicon) for our formal language $\mathfrak{L}$.

2.1.1 Primitive Vocabulary of $\mathfrak{L}$

- The language $\mathfrak{L}$ contains the following primitive expression types.
  - **Connectives**: $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$
  - **Variables**: $x, y, z, \ldots$
  - **Individual Constants**: $a, b, c, \ldots$
  - **1-Place Predicates**: $F^1, F^2, F^3 \ldots F^n$
  - **2-Place Predicates**: $R^1, R^2, R^3 \ldots R^n$
  - ***n*-Place Predicates**: $G_n{}^1, G_n{}^2, G_n{}^3 \ldots G_n{}^n$
  - **Quantifiers**: $\forall, \exists$
  - **Parentheses**: (,)

We will add numerical superscripts to variables and constants when we need more than the alphabet provides.

- Variables and constants are both referred to as *terms*.

2.1.2 Syntax of $\mathfrak{L}$

- Next, we state the syntactic rules of $\mathfrak{L}$, i.e. the rules that determine whether a string of $\mathfrak{L}$ is a well formed formula (wff).

  1. If Π is an n-place predicate and $\alpha_1 \ldots \alpha_n$ are terms, then $\Pi\alpha_1 \ldots \alpha_n$ is wellformed formula (wff).
  2. If ϕ and ψ are wffs, and α is a variable, then the following are wffs:

$$\neg\phi \quad | \quad \phi \wedge \psi \quad | \quad \phi \vee \psi \quad | \quad \phi \rightarrow \psi \quad | \quad \phi \leftrightarrow \psi \quad | \quad \forall\alpha\phi \quad | \quad \exists\alpha\phi$$

3. Only strings formed on the basis of 1. and 2. are wffs.

2.1.3 Variable Binding

- A formula is *closed* if and only if all of its variables are bound. The notion of binding is defined as follows.

DEFINITION: BINDING
A variable α in a wff ϕ is bound in ϕ iff α is within an occurrence of some wff of the form $\forall\alpha\psi$ or $\exists\alpha\psi$ within ϕ. Otherwise α is free.

- For example, in the fomulas below, x is bound but y is free.

$$\exists x\big(F(x)\big) \quad | \quad \exists x\big(F(y)\big) \quad | \quad \forall x\big(R(x,y)\big) \quad | \quad \forall x\big(F(x) \wedge G(y)\big)$$

- Notice that both open and closed formulas count as wffs
- Finally, notice that $\forall$ and $\exists$ are duals, so:

$$\forall\alpha\phi \longleftrightarrow \neg\exists\alpha\neg\phi \quad | \quad \exists\alpha\phi \longleftrightarrow \neg\forall\alpha\neg\phi$$

2.1.4 Semantics and Models for $\mathfrak{L}$

- Next, we need a semantics for $\mathfrak{L}$, i.e. a compositional method for determining the conditions under which the wffs of $\mathfrak{L}$ are true or false. This requires a model for $\mathfrak{L}$.
- A **model** $\mathfrak{M}$ is an ordered pair $\langle\mathcal{D}, \mathcal{F}\rangle$, where:
  - $\mathcal{D}$ is a non-empty set (the domain).
  - $\mathcal{F}$ is an interpretation function which satisfies the following two conditions:
    1. if α is a constant, then $\mathcal{F}(\alpha) \in \mathcal{D}$.
    2. if Π is an n-place predcate, then $\mathcal{F}(\Pi)$ is an n-place relation over $\mathcal{D}$.
- In short, the model provides an extension (i.e. meaning) of the non-logical constants, viz. individual constants and predcate constants.
- Next, we require a **recursive definition of truth** for the wffs of $\mathfrak{L}$, but since our vocabulary now includes variables and quantifiers, we need to say something about the interpretation of these expressions.

2.1.5 Variables in $\mathfrak{L}$

- A variable assignment g for a model $\mathfrak{M}$ is a function from variables in the object language $\mathfrak{L}$ to objects in the domain $\mathcal{D}$. I.e. let VAR be the set of variables, then:

$$g\colon VAR \longmapsto \mathcal{D}$$

- Here is an example of a variable assignment g:

$$g = \begin{bmatrix} x & \longrightarrow & \text{Bob} \\ y & \longrightarrow & \text{Sue} \\ z & \longrightarrow & \text{Mary} \\ \vdots & & \vdots \end{bmatrix}$$

- Hence, we should distinguish between the interpretation of constants and variables.
- Constants are interpreted relative to the interpretation function $\mathcal{F}$ in $\mathfrak{M}$ whereas variables are interpreted relative to a variable assignment g.
- Let $[\![\alpha]\!]_{\mathfrak{M},g}$ stand for the denotation of α relative to $\mathfrak{M}$ and g. So,

$$[\![\alpha]\!]_{\mathfrak{M},g} = \begin{cases} \mathcal{F}(\alpha) & \text{if } \alpha \text{ is a } \textit{constant} \\ g(\alpha) & \text{if } \alpha \text{ is a } \textit{variable} \end{cases}$$

2.1.6 Valuations and Truth-in-a-Model

- A valuation function $\mathcal{V}$ for a model $\mathfrak{M}$ and some variable assignment g is a function which assigns to each wff either 0 or 1 under the following constraints.
  - For any n-place predcate Π and any terms $\alpha_1...\alpha_n$, $\mathcal{V}_{\mathfrak{M},g}(\Pi\alpha_1...\alpha_n) = 1$ iff $\langle [\alpha_1]_{\mathfrak{M},g}...[\alpha_n]_{\mathfrak{M},g}\rangle \in \mathcal{F}(\Pi)$
  - For any wffs ϕ, ψ, and any variable α:

$$\begin{array}{rll}
\mathcal{V}_{\mathfrak{M},g}(\neg\phi) & = 1 \text{ iff} & \mathcal{V}_{\mathfrak{M},g}(\phi) = 0 \\
\mathcal{V}_{\mathfrak{M},g}(\phi \wedge \psi) & = 1 \text{ iff} & \mathcal{V}_{\mathfrak{M},g}(\phi) = 1 \textbf{ and } \mathcal{V}_{\mathfrak{M},g}(\psi) = 1 \\
\mathcal{V}_{\mathfrak{M},g}(\phi \vee \psi) & = 1 \text{ iff} & \mathcal{V}_{\mathfrak{M},g}(\phi) = 1 \textbf{ or } \mathcal{V}_{\mathfrak{M},g}(\psi) = 1 \\
\mathcal{V}_{\mathfrak{M},g}(\phi \rightarrow \psi) & = 1 \text{ iff} & \mathcal{V}_{\mathfrak{M},g}(\phi) = 0 \textbf{ or } \mathcal{V}_{\mathfrak{M},g}(\psi) = 1 \\
\mathcal{V}_{\mathfrak{M},g}(\forall\alpha\phi) & = 1 \text{ iff} & \textbf{for every } u \in \mathcal{D}, \mathcal{V}_{\mathfrak{M},g[u/\alpha]}(\phi) = 1 \\
\mathcal{V}_{\mathfrak{M},g}(\exists\alpha\phi) & = 1 \text{ iff} & \textbf{for at least one } u \in \mathcal{D}, \mathcal{V}_{\mathfrak{M},g[u/\alpha]}(\phi) = 1
\end{array}$$

- We now define truth-in-a-model as follows.

> **DEFINITION:** TRUTH-IN-A-MODEL
> ϕ is *true in* a model $\mathfrak{M}$ iff $\mathcal{V}_{\mathfrak{M},g}(\phi) = 1$, for each variable assignment g for $\mathfrak{M}$.

2.1.7 Validity and Logical Consequence

- Validity is defined as truth in all models.

> **DEFINITION:** VALIDITY
> ϕ is valid in $\mathfrak{L}$ iff ϕ is true in all models $\mathfrak{M}$. $\vDash_{\mathfrak{L}} \phi$

- Logical consequence is defined in terms of truth preservation.

> **DEFINITION:** LOGICAL CONSEQUENCE
> ϕ is a logical consequence of a set of wffs Γ in $\mathfrak{L}$ iff:
> For every model $\mathfrak{M}$ and every variable assignment g for $\mathfrak{M}$, if $\mathcal{V}_{\mathfrak{M},g}(\gamma) = 1$ for every $\gamma \in \Gamma$, then $\mathcal{V}_{\mathfrak{M},g}(\phi) = 1$. $\Gamma \vDash_{\mathfrak{L}} \phi$

Chapter 3

Set Theory

3.1 Sets, Relations, and Functions

More information than that contained in these notes can be found in e.g. Partee *et al.* (1990).

- A set is an abstract collection of elements which are the *members* of the set. Anything can be a member of a set, e.g. other sets. In the examples in these notes I will mostly use sets whose members are natural numbers.

Set membership: $\in$

- We indicate that an element x is a member of a set A is the symbol '$\in$'. I.e.

$$x \in A$$

- There are many ways to specify the members of a set. One is **list notation** which is useful when the members can easily be written down.
  - {1, 2, 3, 17, Barack Obama, the king of Sweden}

- A set can also be specified using so-called **predicate notation** where the predicate states one (or more) condition(s) that is a necessary and sufficient condition for being a member of the set.
  - $\{x : x$ is natural number$\}$
  - $\{y : y$ was born in Germany$\}$
  - $\{x \mid x$ is natural number$\}$
  - $\{y \mid y$ was born in Germany$\}$

- A third way of specifying a set is by using (possibly recursive) rules, i.e.

$$\left.\begin{array}{ll} 1. & 0 \in A. \\ 2. & \text{if } x \in A, \text{ then } x + 1 \in A. \\ 3. & \text{Nothing else is a member of } A. \end{array}\right\} \{x : x \text{ is a natural number}\}$$

- The set that has no members is called **the empty set** (or sometimes *the null set*). This set is normally denoted by '$\varnothing$'. A set that has exactly one member is called **a singleton set**.

Singleton sets and the empty set: $\varnothing$

Cardinality

Cardinality: $|A|$

- If considering the number of elements of a set A, we are considering the cardinality of A. For example, the cardinality of the set {5, Barack Obama, The Eiffel Tower} is 3.
- We indicate the cardinality of a set A by enclosing it in vertical lines as follows: $|A| = 3$.

3.1.1 Subsets and Power Sets

- Subsets are defined as follows.

> **DEFINITION:** SUBSET
> A is a **subset** of another set B, $A \subseteq B$, iff every element of A is a member of B.

But notice also that the claims: $A \subseteq A$ and $A \in A$ are not equivalent.

Proper Subset: $A \subset B$

  - Note, given the definition of subset above, every set is a subset of itself. That is, for every set A, $A \subseteq A$.
  - In cases where a set B has more members than a subset A, we say that A is a **proper subset** of B, and we write that as follows: $A \subset B$.

- Powersets are defined as follows.

> **DEFINITION:** POWERSET
> The powerset $\mathcal{P}$ of a set A is the set of subsets of A: $\mathcal{P}(A) \overset{\text{def}}{=} \{B : B \subseteq A\}$

  - If $A = \{1, 2, 3\}$, then, $\mathcal{P}(A) = \{\varnothing, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$
  - Notice that the power set of some set A always includes the empty set. In other words, for any set A, it is always the case that $\varnothing \subseteq A$.

3.1.2 Set Theoretic Operations: Union, Intersection, Complementation

Union: $A \cup B$

- Unions are defined as follows.

> **DEFINITION:** UNION
> The **union** of two sets A and B is the set of elements that are members of A **or** B:
> $A \cup B \overset{\text{def}}{=} \{x : x \in A \text{ or } x \in B\}$

  - In the Venn diagram below, the union of A and B is the set of elements contained in the area colored green, i.e. every element that is either in A or in B.

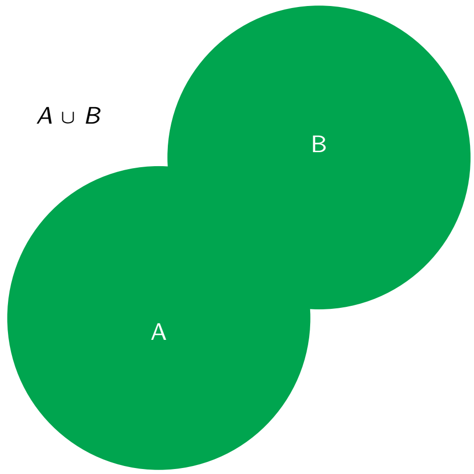

Intersection: A∩B

- Intersection is defined as follows.

> **DEFINITION**: INTERSECTION
> The intersection of two sets A and B is the set of elements that are members of *both* A and B: $A \cap B \stackrel{\text{def}}{=} \{x : x \in A \text{ and } x \in B\}$

- In the Venn diagram below, the intersection of A and B is the set of elements contained in the area colored green, i.e. every element that is both in A and in B.

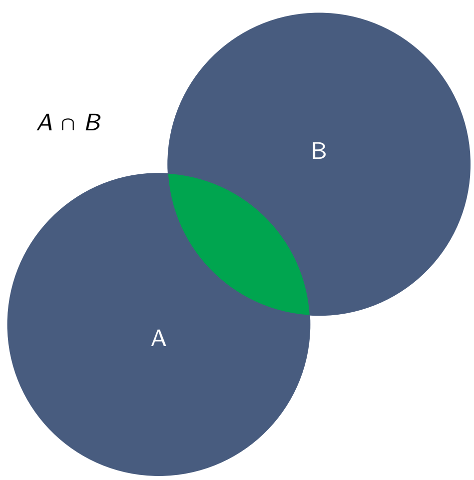

Relative Complement: A–B, also written $A \setminus B$

▸ Relative complement of A in B is defined as follows.

DEFINITION: RELATIVE COMPLEMENT
The relative complement of B in A, A–B, is the set of elements that are members of A but not members of B, i.e. it is A with the members of B "subtracted": $A–B \stackrel{\text{def}}{=} \{x : x \in A \text{ and } x \notin B\}$

· In the Venn diagram below, the complement of B in A is the set of elements contained in the area colored green, i.e. the set of elements that are in A but not in B.

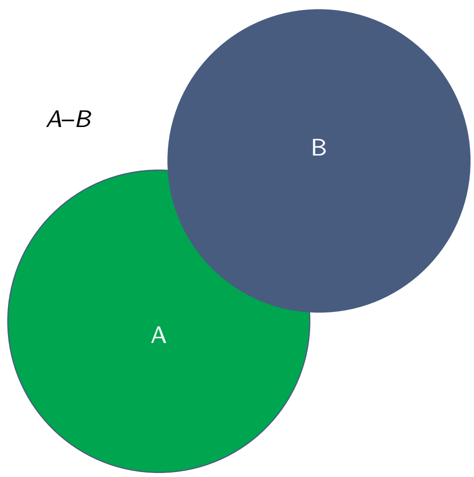

▸ The complement (simpliciter) of A is defined as follows:

DEFINITION: COMPLEMENT
The complement of a set A, A', is all the elements of the universe U that are not members of A: $A' \stackrel{\text{def}}{=} \{x : x \notin A\}$.

3.1.3 Set-Theoretic Equivalencies

IDEMPOTENCY
$X \cup X = X$
$X \cap X = X$

COMMUTATIVITY
$X \cup Y = Y \cup X$
$X \cap Y = Y \cap X$

DEMORGAN'S LAW
$(X \cup Y)' = X' \cap Y'$
$(X \cap Y)' = X' \cup Y'$

COMPLEMENTS
$X \cup X' = U$
$(X')' = X$

ASSOCIATIVITY
$(X \cup Y) \cup Z = X \cup (Y \cup Z)$
$(X \cap Y) \cap Z = X \cap (Y \cap Z)$

IDENTITIES
- $X \cup \varnothing = X$
- $X \cap \varnothing = \varnothing$

3.2 Relations and Functions

3.2.1 Ordered Sets

- The sets considered above are unordered — in particular, the following identity relation holds:

$$\{1,2\} = \{2,1\}$$

- In order to use set theory to model relations among objects, we need sets that are ordered, viz. ordered sets.
- However, we can define an ordered set (an ordered pair) in terms of an unordered set as follows: Ordered Pairs

$$\langle 1,2 \rangle \stackrel{\text{def}}{=} \{\{1\},\{1,2\}\}$$

- Now, $\langle 1,2 \rangle$ denotes the set whose first member is 1, and whose second member is 2. Given this definition, it now follows that $\langle 1,2 \rangle \neq \langle 2,1 \rangle$ since:

$$\{\{1\},\{1,2\}\} \neq \{\{2\},\{1,2\}\}$$

- To make things easier for ourselves, we will pretty much treat ordered sets as primitives from here on out.

3.2.2 (Cartesian) Products

- From sets A and B, we can construct a set of ordered pairs where the first element of each pair is a member of A and the second element is a member of B — this is called the Cartesian product of A and B and it is written $A \times B$.

DEFINITION: CARTESIAN PRODUCT
The (cartesian) product of A and B is the set of pairs $\langle a,b\rangle$ where $a \in A$ and $b \in B$:
$A{\times}B \stackrel{\text{def}}{=} \{\langle x,y\rangle : x \in A \text{ and } y \in B\}$

3.2.3 Relations

Relations

- Objects are often related in various ways. For example, the 'father of'-relation denotes a relation between a man and his offspring.
- To indicate that an object a bears the R-relation to another object b (possibly a itself), we write Rab — or alternatively aRb or $R(a,b)$.
- To indicate that a relation R holds between objects of two sets A and B, we write $R \subseteq A{\times}B$.
- For example, if R is the relation 'father of', then R is a relation that holds between objects of the set A and the set B as below.

$$A = \{18, \text{Jack}, \text{Sam}\} \qquad B = \{\text{Sam Jr.}, \text{Jack Jr.}\}$$

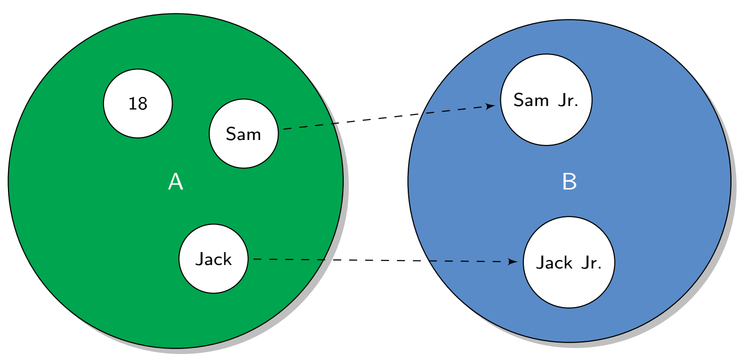

- The cartesian product of A and B is as follows.

$$\begin{array}{lcll} A{\times}B & = & \{ & \langle 18, \text{Sam Jr.}\rangle, \langle 18, \text{Jack Jr.}\rangle, \langle \text{Sam}, \text{Sam Jr.}\rangle \\ & & & \langle \text{Sam}, \text{Jack Jr.}\rangle, \langle \text{Jack}, \text{Sam Jr.}\rangle, \langle \text{Jack}, \text{Jack Jr.}\rangle \quad \} \end{array}$$

- Jack is the parent of Jack Jr. — so, $\langle$Jack, Jack Jr.$\rangle \in R$ and Sam is the parent of Sam Jr. — so, $\langle$Sam,Sam Jr.$\rangle \in R$.
  - Hence $R = \{\langle\text{Jack,Jack Jr.}\rangle, \langle\text{Sam,Sam Jr.}\rangle\}$ and $\{\langle\text{Jack,Jack Jr.}\rangle, \langle\text{Sam,Sam Jr.}\rangle\} \subseteq A{\times}B$.
  - Notice, however, that R (when meaning 'father of') is not a member of $B{\times}A$.

Complement of a Relation

- The complement of a relation R is defined as follows.

DEFINITION: COMPLEMENT OF RELATION
The complement of a relation is the set of ordered pairs that are not members of R: $R' \stackrel{\text{def}}{=} (A \times B) — R$

· For example, if R is the relation 'father of' (as above), we would have the following:

$$R' = \{\langle 18,\text{Jack Jr.}\rangle, \langle 18, \text{Sam Jr.}\rangle, \langle \text{Jack},\text{Sam Jr.}\rangle, \langle \text{Sam},\text{Jack Jr.}\rangle\}$$

- The inverse of a relation R is defined as follows.

The Inverse of a Relation

DEFINITION: INVERSE OF RELATION
The inverse of a relation $R \subseteq A \times B$ has as its members all the ordered pairs in R but with the first and second coordinate of each ordered pair reversed. Therefore, where $R \subseteq A \times B$: $R^{-1} \stackrel{\text{def}}{=} \{\langle y,x\rangle \in B \times A \mid \langle x,y\rangle \in R\}$

· So, if R is the relation 'father of' (as above), we would have the following:

$$R^{-1} = \{\langle \text{Jack Jr.}, \text{Jack}\rangle, \langle \text{Sam Jr.},\text{Sam}\rangle \}$$

Terminology: Domains and Ranges of Relations

Domains and Ranges

- If $R \subseteq A \times B$, then the set containing the first coordinates of the ordered pairs in R equals the domain of R. The set containing the second coordinates of the ordered pairs in R equals the range of R.

3.2.4 Functions

- We are now ready to consider one of the most central concepts in mathematics and logic; a concept absolutely crucial to formal semantics — the concept of a **function**.

Functions

- A function is a special kind of relation between sets — in particular a relation that satisfies certain conditions.
- We define the concept of a function as follows.

DEFINITION: FUNCTION
A relation $R \subseteq A \times B$ is a function if and only if:

1. Each element in the domain of R is paired with *at most* one element in the range of R.

We will say that R is a **total** function if R satisfies both (1.) above and (2.) below. If it satisfies only one, R is a **partial** function.

2. *Every* element of A is mapped to some element in B.

- Let's consider a couple of examples. Suppose we have the following sets:

$$A = \{a,b,c,d\} \qquad B = \{1,2,3\}$$

A relation $R \subset A \times B$ which is a total function f
f: $A \longmapsto B$

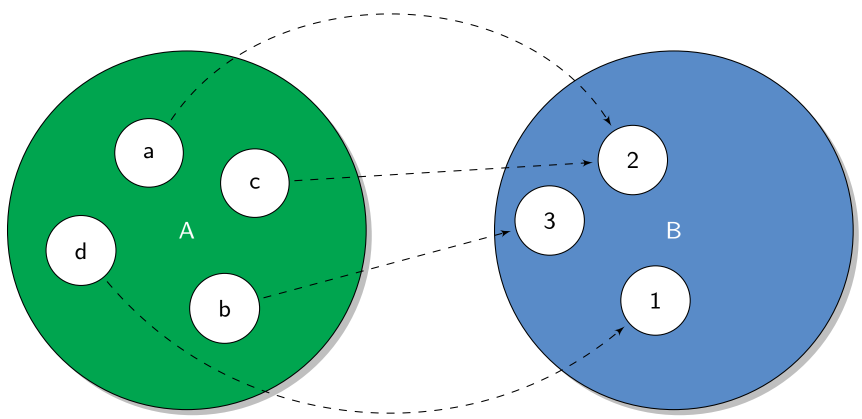

$\{\langle a,2\rangle, \langle b,3\rangle, \langle c,2\rangle, \langle d,1\rangle\}$

Another relation $R \subset A \times B$ which is a total function g
g: $A \longmapsto B$

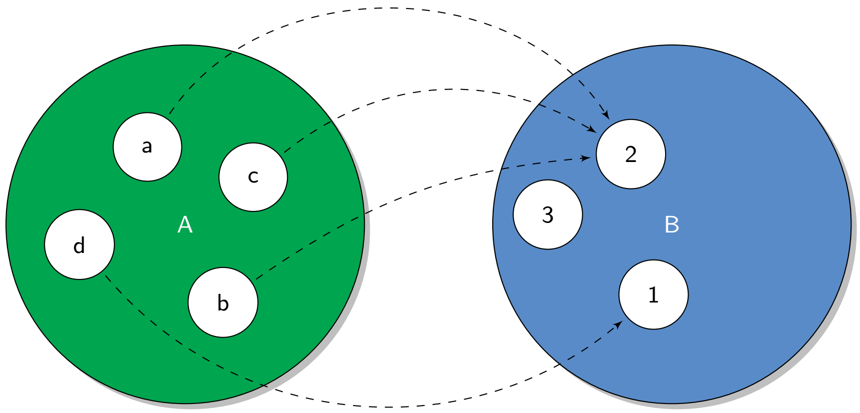

$\{\langle a,2 \rangle, \langle b,2 \rangle, \langle c,2\rangle, \langle d,1\rangle\}$

- **Some Terminology**
  - If $\langle a,b\rangle$ is an element of a relation R and R is a function, a is sometimes referred to as the *input* or *argument* of the function, and b is often referred to as the *output* or *value* of the function.
  - Functions are often also referred to as a *mapping* (or simply a *map*) from A to B.
  - Informally, you can think of functions as machines that spit out values y (from the range of the function) when fed an argument x (from the domain of the function).

A relation $R \subset A \times B$ which is a partial function

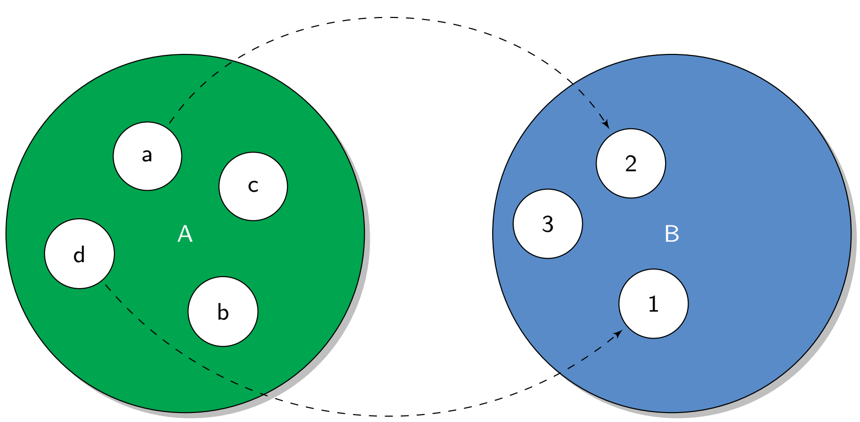

{⟨a,2⟩, ⟨d,1⟩}

- A simple example of a partial function is *division* (÷) in arithmetic. Division is a function from ordered pairs of numbers (argument) to numbers (value) — but the function is *not defined* when the divisor is 0.
- Hence, while the domain of the function is the product of the set of, say, natural numbers ($\mathbb{N}\times\mathbb{N}$), the function is not defined for the elements of that set where the second coordinate is 0.
- So, *division* fails to satisfy (2) and is hence a partial function.

▸ Another way to think of a partial function is simply as a total function on a subset of the relevant domain.

A relation $R \subset A\times B$ which is not a function

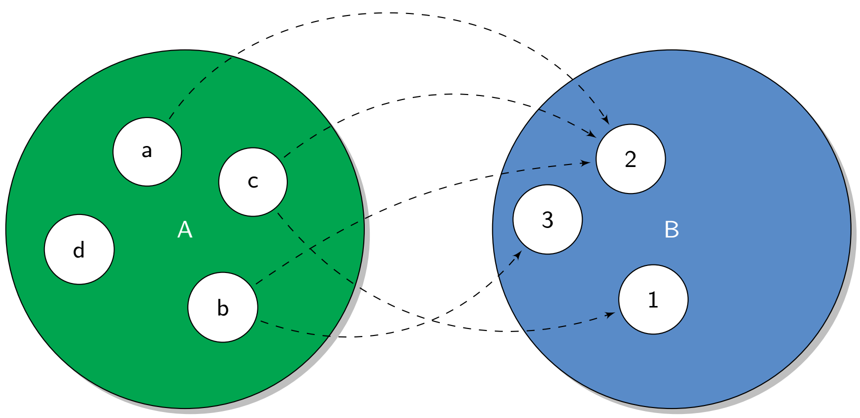

{⟨a,2⟩, ⟨b,2 ⟩, ⟨b,3⟩, ⟨c,1⟩, ⟨c,2⟩}

Homework: Explain why is this not a function.

Properties of Functions

▸ Functions that never map two elements of the domain A to the same element of the range B are called **injective** (also called **one-to-one functions** or functions from A **into** B).

▸ Functions where for every element in the range B there is a mapping from one (or more) elements in the domain A are said to be **surjective** (also called functions from A **onto** B).

▸ And, finally, functions that are both injective and surjective, i.e. for every element of the range B, there is exactly one mapping from a distinct element of A, are said to be *bijective* (also called **one-to-one correspondences**).

▸ Bijective functions are special because their inverses are always total functions.

Injective **Surjective** **Bijective**

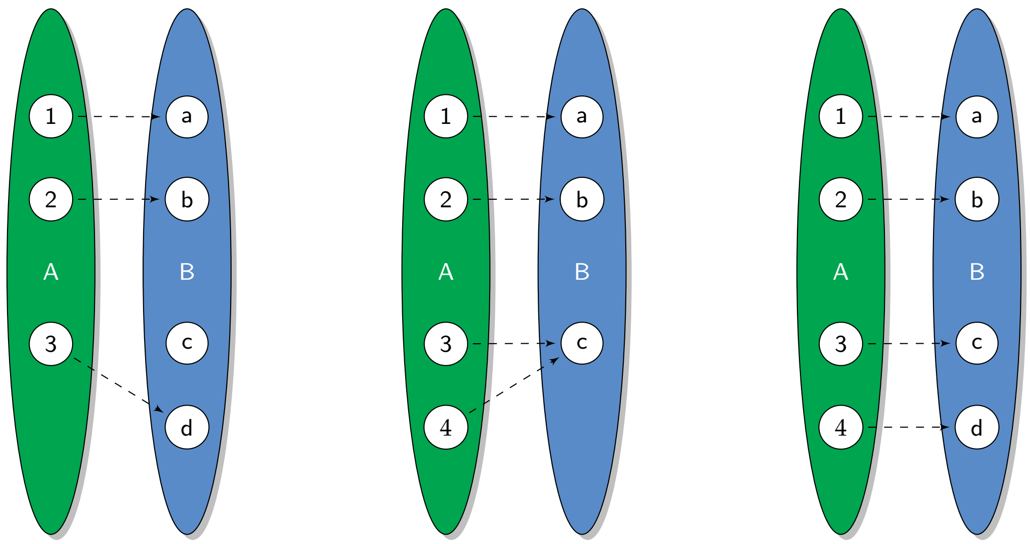

Function Composition

- From two functions $f\colon A \longmapsto B$ and $g\colon B \longmapsto C$, we can compose a new function $h\colon A \longmapsto C$. This is called the composition of f and g.
- The composition of f and g is indicated using the symbol ⌜$\circ$⌝, viz. $g \circ f$.
- Function composition is defined as follows.

NB! The order might seem backwards here, but there is an important reason for that: f is the first function, and g is the second. The value of f at an argument a is $f(a)$, and the value of g at the argument $f(a)$ is written $g(f(a))$.

DEFINITION: FUNCTION COMPOSITION
The composition of f and g is the set of pairs $\langle x,z \rangle$ where for some element y, y is such that $\langle x,y \rangle$ is a member of f and $\langle y,z \rangle$ is a member of g:

$g \circ f \overset{\text{def}}{=} \{\langle x,z \rangle\text{: for some } y,\ \langle x,y \rangle \in f \text{ and } \langle y,z \rangle \in g\}$

NB! $f \circ g$ is not generally equal to $g \circ f$.

$\boldsymbol{f}$**:** $\boldsymbol{A \longmapsto B}$

$f = \{\langle 1,\text{a}\rangle, \langle 2,\text{b}\rangle, \langle 3,\text{b}\rangle\}$

$\boldsymbol{g}$**:** $\boldsymbol{B \longmapsto C}$

$g = \{\langle \text{a},6\rangle, \langle \text{b},7\rangle, \langle \text{c},7\rangle\}$

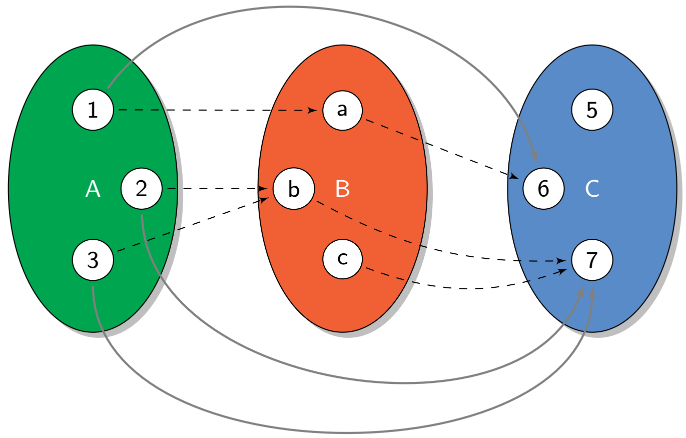

$\boldsymbol{g \circ f = h}$**:** $\boldsymbol{A \longmapsto C}$

$h = \{\langle 1,6\rangle, \langle 2,7\rangle, \langle 3,7\rangle\}$

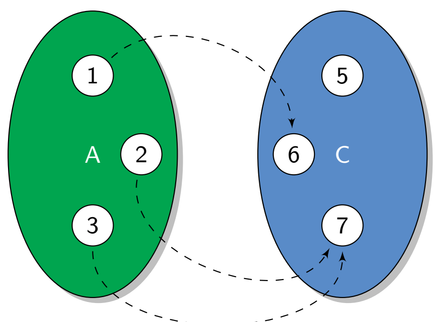

- When trying to compose two functions it might happen that the range of the first function contains no elements in the domain of the second. In that case, the result of composing the functions is the empty set.

Composing Relations

- Not only functions can be composed—we can also compose relations.
- For example, suppose we have to relations $R \subseteq A \times B$ and $S \subseteq B \times C$. If so, then...

$$S \circ R = \{\langle x,z\rangle\text{: for some } y,\ \langle x,y\rangle \in R \text{ and } \langle y,z\rangle \in S\}$$

Chapter 4

Formal Semantics

4.1 Compositional Derivations of Truth Conditions

- Frege proposed a fundamental distinction between **saturated** and **unsaturated** meanings — he considered this distinction to be essential to the compositional nature of meaning. For example, consider (27).

  (27) Sue snores.

- This sentence can be split into two parts, namely 'Sue' and 'snores'.
- Frege's conjecture is that the meaning of (27) can be derived by applying the **unsaturated** part of the sentence to the **saturated** part. You can think of saturated and unsaturated parts as follows.
  - Unsaturated parts are **functions** (functions that take arguments and output values)
  - Saturated parts are **individuals/object** from the domain (arguments for functions)
- Thus, the general Fregean idea of meaning yields the following crude picture of the compositional meaning of (27).

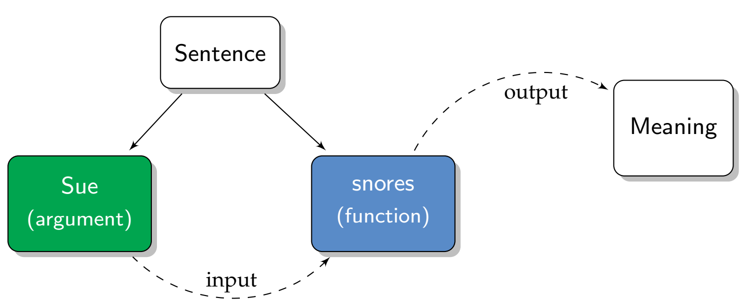

- On the picture of natural language meaning that we adopt here, the meaning of a natural language sentence S must thus be explicated in terms of only three factors.
  - The meaning of the parts of S.
  - The syntactic structure of S (the way the parts are combined).
  - Function application, i.e. pairwise applications of functions to arguments.
- The general aim in this course is to explore to what extent such an analysis is feasible and what obstacles it faces.

4.1.1 Three Main Obstacles

- Since our main aim is to provide a recursive procedure for deriving the truth conditions of sentences by only considering the meaning of the individual parts and their syntactic mode of combination, there are three immediate questions that need to be addressed.
  - FIRST OBSTACLE
    We need to know what basic expressions mean (or denote). For example, for (27), we need to know what the denotation of e.g. a **proper name** is and what the denotation of an **intransitive verb** is.
  - SECOND OBSTACLE
    We need to state the rules for deriving the meaning of a complex expression (e.g. a sentence) from the meanings of the basic expressions. Function application is just one example of such a rule (but we might need additional rules).
  - THIRD OBSTACLE
    We need a way of determining the structural relations obtaining within a sentence. For example, consider (28).

    (28) A barber from Seville lost all his money on the black market.

    Which parts should we split this sentence into? I.e. which expressions are supposed to combine pairwise with which expressions?

4.1.2 First Obstacle: Denotational Domains and Interpretation Functions

- The meaning of both basic and complex expressions (within a formal semantic theory) are standardly referred to using names such as, e.g. *extension*, *denotation*, and *semantic value*. I use these interchangebly in these notes.
- As in FOL, giving a semantics for a language (in this case English) requires a model. We will assume that our model $\mathfrak{M}$ is an ordered tuple consisting of a domain of individuals D_e, a domain of truth values D_t, and an interpretation function $\mathcal{I}$.

$$\mathfrak{M} = \langle D_e, D_t, \mathcal{I} \rangle$$

1. **Domain $\mathbf{D}_e$**
   D_e is the set of individuals/objects in the universe under consideration.

2. **Domain $\mathbf{D}_t$**
   D_t is the set of truth values, {0,1}.
3. **Interpretation Function $\mathcal{I}$**
   The interpretation function is a function from basic expressions of the language (e.g. proper names, predicates etc.) to their extensions.

- We will assume that the interpretation function assigns the following denotations to basic expressions. For example:

  NB! When talking explicitly about expressions of the **object language**, the expression will be in italics.

  · $\mathcal{I}$ maps **proper names** to **individuals** (objects in D_e):

  $$\mathcal{I}(Sue) = \text{Sue}$$

  · $\mathcal{I}$ maps **intransitive verbs** to **functions** from D_e to D_t (subsets of $D_e{\times}D_t$).

$$\mathcal{I}(snore) = \Big[\ f\colon D_e \longmapsto \{0,1\} \qquad \text{for all } x\text{: } f(x) = \begin{cases} 1 \text{ if } x \in \{x \in D_e \mid x \text{ snores}\} \\ 0 \text{ if } x \notin \{x \in D_e \mid x \text{ snores}\} \end{cases} \Big]$$

Characteristic Functions

- Notice that verbs are here assumed to denote functions. This is a minor (but innocuous) departure from FOL, where verbs (i.e. predicates) denote sets of individuals.
- Nevertheless, it is for various reasons convenient to be able to speak of intransitive verbs as if they simply denoted sets of individuals, e.g. the following set.

$$\mathcal{I}(snore) = \{x \in D_e \mid x \text{ snores}\}$$

- However, from such a set, we can define a so-called characteristic function. Suppose A is the set below — and notice that $A \subseteq D_e$

$$A = \{x \in D_e \mid x \text{ snores}\}$$

- The characteristic function of A (over D_e) is the function f^A: $D_e \longmapsto D_t$ which is defined as follows:

$$\text{for all } x \in D_e\text{:} \quad f^A(x) = \begin{cases} 1 \text{ if } x \in A \\ 0 \text{ if } x \notin A \end{cases}$$

- That is, the characteristic function of A (over D_e) is the function f^A such that:

$$\text{for all } x \in D_e\text{:} \quad f^A(x) = \begin{cases} 1 \text{ if } x \in \{x \in D_e \mid x \text{ snores}\} \\ 0 \text{ if } x \notin \{x \in D_e \mid x \text{ snores}\} \end{cases}$$

- So, if D_e = {Sue, Kim, Ted, Bob} and only Sue and Kim snore, viz. A = {Sue, Kim}, then the characteristic function of A (over D_e) is the function f^A illustrated below.

$\boldsymbol{f^A}$**:** $\boldsymbol{D_e \longmapsto D_t}$

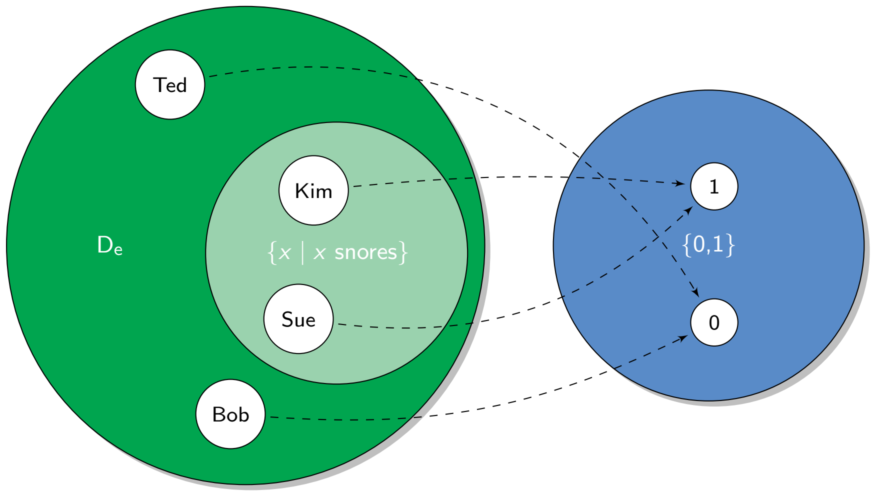

{⟨Sue,1⟩, ⟨Kim,1 ⟩, ⟨Bob,0⟩, ⟨Ted,0⟩}

Sets and Characteristic Functions: Interdefinability

- It is important to notice the following.
  1. For any function $f\colon D_e \longmapsto D_t$, we can identify a subset A of D_e of which that function is the characteristic function (e.g. the set of individuals who snore.)
  2. Given any subset A of D_e, we can identify A's characteristic function over D_e.
  3. So, there is a one-to-one correspondence between sets and characteristic functions.
- Hence, speaking informally as if e.g. intransitive verbs denote sets of individuals is harmless. When convenient, we will therefore sometimes use set talk rather than function talk.

4.1.3 Second Obstacle: The Denotation of Complex Expressions

- Our interpretation function $\mathcal{I}$ in our model $\mathfrak{M}$ only determines the meaning of basic expressions, but we need a procedure for determining the meaning of complex expressions, e.g. (27)

  (27) Sue snores.

- We therefore introduce a second generalized interpretation function $[\![\cdot]\!]$ — the equivalent of a valuation function in FOL.
- Where $\mathfrak{M} = \langle D_e, D_t, \mathcal{I}\rangle$, we assume the following for basic expressions θ of the language:

$$[\![\theta]\!]^{\mathfrak{M}} = \mathcal{I}_{\mathfrak{M}}(\theta)$$

From here on out, I generally omit the superscript $\mathfrak{M}$ on the interpretation function. This is a bit sloppy, but just keep in mind that the interpretation function must always be relativized to a model.

- Next, we generalize our new interpretation function $[\![\cdot]\!]$ to complex expressions by introducing composition rules. For example.

  - If an expression α is of the form: $\begin{array}{c} S \\ \beta \quad\quad \gamma \end{array}$ then $[\![\alpha]\!] = [\![\gamma]\!]([\![\beta]\!])$

  - Hence, $\left[\!\!\left[\begin{array}{c} S \\ \textit{Sue} \quad\quad \textit{snores} \end{array} \right]\!\!\right] = [\![\textit{snores}]\!]([\![\textit{Sue}]\!])$

- This rule states that the extension of a sentence such as (27) is determined by applying the function denoted by 'snores' to the argument denoted by 'Sue'.

4.1.4 Deriving the Meaning of 'Sue snores'

- This is sufficient to derive the meaning, i.e. the truth conditions, of (27) — as demonstrated below. Our derivation should yield the result that the sentence 'Sue snores' = 1 if and only if Sue snores.

1. 'Sue snores' = 1 iff Sue snores.

2. 'Sue snores' has the form: $\begin{array}{c} \text{S} \\ \text{Sue} \quad\quad \text{snores} \end{array}$

3. From our composition rule it follows: $\left[\!\!\left[\begin{array}{c} S \\ \textit{Sue} \quad\quad \textit{snores} \end{array} \right]\!\!\right] = [\![\textit{snores}]\!]([\![\textit{Sue}]\!])$.

4. By transitivity of identity: $[\![\textit{snores}]\!]([\![\textit{Sue}]\!]) = 1$ iff Sue snores

5. By the definition of our interpretation function: $[\![\textit{Sue}]\!] = \mathcal{I}(\textit{Sue})$ = Sue.

6. Hence, $[\![\textit{snores}]\!]$(Sue) = 1 iff Sue snores

7. By the definition of our interpretation function:

$$[\![\textit{snores}]\!] = \mathcal{I}(\textit{snores}) = \left[\, f\colon D \longmapsto \{0,1\} \quad \text{for all } x\text{: } \begin{array}{l} f(x) = 1 \text{ if } x \in \{x \mid x \text{ snores}\} \\ f(x) = 0 \text{ if } x \notin \{x \mid x \text{ snores}\} \end{array} \right]$$

8. Hence,

$$\left[\, f\colon D \longmapsto \{0,1\} \quad \text{for all } x\text{: } \begin{array}{l} f(x) = 1 \text{ if } x \in \{x \mid x \text{ snores}\} \\ f(x) = 0 \text{ if } x \notin \{x \mid x \text{ snores}\} \end{array} \right](\text{Sue}) = 1 \text{ iff Sue snores}$$

- We have succeeded. Relying on only the meaning of the parts of (27), the syntactic structure of (27), and our rule of function application, we have derived what appears to be the correct truth conditions for (27).
- However, (27) is a very simple sentence and if this strategy is to succeed more generally, we will need to complicate our theory quite a bit. For now we turn to the third obstacle.

4.2 Syntax and Semantics: Phrase Structures as Input to Semantics

We did not define $[\![\cdot]\!]$ for sentences, but for structures.

- If you were paying close attention, you would have noticed that we have not actually derived the truth conditions for a sentence of English, but rather for a kind of complex structure.
- That is, our interpretation function $[\![\cdot]\!]$ was defined for a structure rather than a sentence.

~~⟦Sue snores⟧~~ 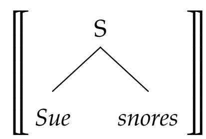

- This might seem strange but there is a good reason for this decision. Most syntacticians today agree that the structure of natural language sentences is much richer than what is visible on the surface — i.e. that sentences are structured entities composed of a number of different syntactic constituents.
- These structures are standardly represented using so-called phrase structure trees. Phrase structure trees make the constituent structure of sentences clear — for example:

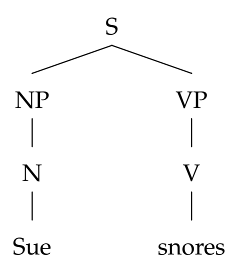

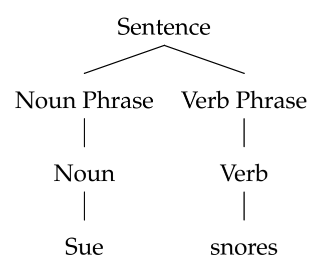

- Within so-called transformational grammars, one often distinguishes between several layers of linguistic representation, e.g.
  - The surface structure of a sentence which is a phrase structure where crudely speaking the words are ordered as they are seen and heard.
  - The logical form (LF) of sentence which is an additional layer of linguistic representation derived using various transformational rules. At LF, the words may not be ordered as they are seen and heard. Later in this course, we will consider a number of arguments for the existence of this extra layer of linguistic representation.

- Many approaches to formal semantics — sometimes referred to as LF-based semantics — take LFs to be the inputs to the semantic machinery (and not simply sentences of English or surface structures).
- This is not a trivial assumption by any means, but it does make life as a semanticist easier in many ways. For example:
  - LFs make the *structure* of various sentence totally transparent, i.e. how lexical items are supposed to combine.
  - Certain ambiguities look very much like *structural* ambiguities — that is, differences in meaning that arise as a result of different constituent structures.

4.2.1 Phrase Structures and Ambiguity

- The sentence in (29) below has two possible meanings, paraphrased in (29a) and (29b).

(29) Sue saw the boy with the binoculars.
  a. Sue saw the boy using a pair of binoculars.
  b. Sue saw the boy who had a pair of binoculars.

- In a transformational grammar, e.g. Government and Binding theory, two different phrase structures are associated with (29), one corresponding to (29a) and one to (29b).
- The prediction that (29) has two possible phrase structures explains why it is ambiguous.

The boxed NPs indicate the object that Sue saw.

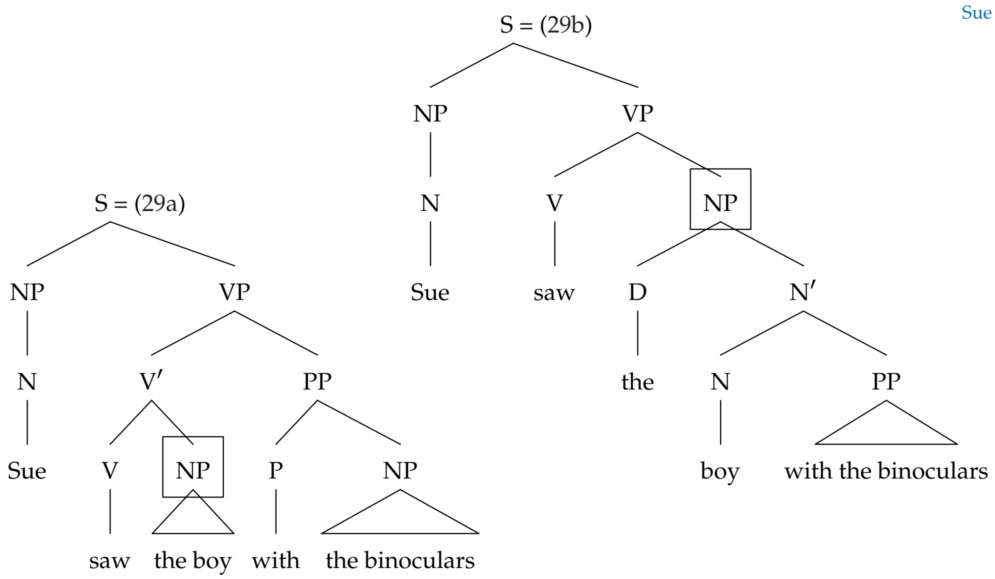

- We will basically assume that the inputs for semantic composition, i.e. what we are trying to derive the truth conditions of, are LFs rather than surface structures.

- This means that syntax alone determines the order in which lexical items combine. Since phrase structures are assumed to be binary branching, this fits quite well with the Fregean idea that compositionality consists in the stepwise application of functions to single arguments.

4.2.2 More Composition Rules

- If phrase structures are the input to semantic composition, we need more than just one composition rule to derive truth conditions. In particular, we need rules to make our interpretation function defined for phrase structures such as (30).

(30) 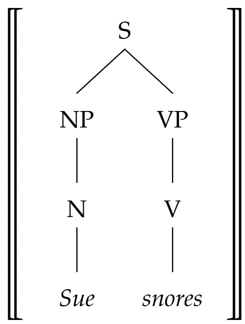 = 1 iff Sue snores

- So in addition to rules determining the denotations of 'Sue' and 'snores' (and their combination), we need rules for the syntactic constituents, e.g. NP, VP, N, V.
- We already have the composition rule S_1.

RULE S_1

If α has the form 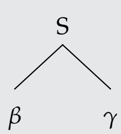 then $[\![\alpha]\!] = [\![\gamma]\!]([\![\beta]\!])$.

- Here are the additional rules we now need.

RULE S_2

If α has the form  then $[\![\alpha]\!] = [\![\beta]\!]$

RULE S3

If α has the form $\begin{array}{c}\text{VP}\\ |\\ \beta\end{array}$ then $[\![\alpha]\!] = [\![\beta]\!]$

RULE S4

If α has the form $\begin{array}{c}\text{N}\\ |\\ \beta\end{array}$ then $[\![\alpha]\!] = [\![\beta]\!]$

RULE S5

If α has the form $\begin{array}{c}\text{V}\\ |\\ \beta\end{array}$ then $[\![\alpha]\!] = [\![\beta]\!]$

- We are starting to proliferate composition rules (which we should try to avoid).
- However, as I hope is obvious, there is a way of greatly simplifying these rules—we'll get to that later.

4.3 Extending the Theory

- Consider the sentence below.

(31) Mia likes Bob.

(32)

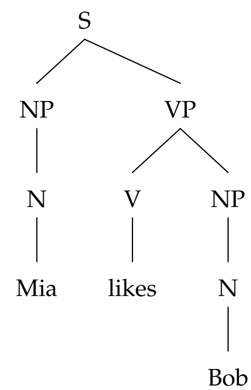

- The phrase structured associated with (31) is given by (32) on the right.
- The question to be addressed in order to derive the truth conditions of (31) is what the denotation of the transitive verb 'like' should be.
- Applying (S1) to (32) yields:

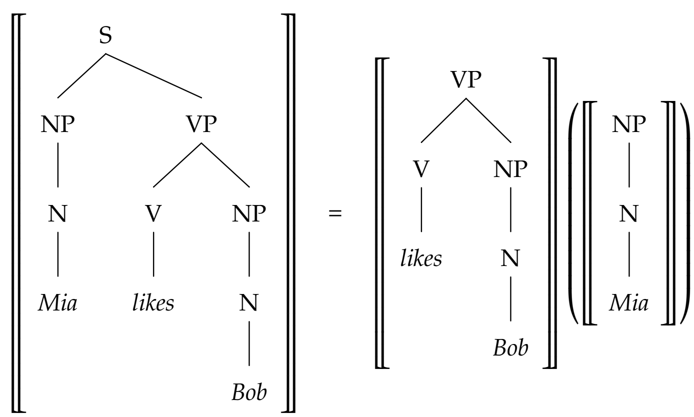

- Applying (S2) and (S4), we can simplify this to:

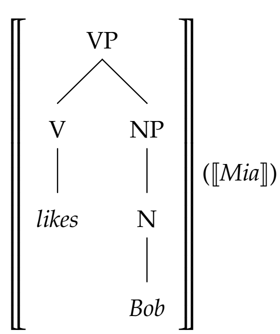

- However, we currently have no rules for branching VP nodes.
- To determine which rule, let's look at the VP-constituent. We know that:
  - $\begin{array}{c}\text{NP}\\ |\\ \text{N}\end{array}$ reduces to $[\![\text{N}]\!]$
  - $\begin{array}{c}\text{N}\\ |\\ \textit{Bob}\end{array}$ reduces to $[\![\textit{Bob}]\!]$.
- $[\![\textit{Bob}]\!]$ = Bob. Since this is not a function, $[\![\textit{Bob}]\!]$ must be an argument for a function.
  - $\begin{array}{c}\text{V}\\ |\\ \text{likes}\end{array}$ reduces to $[\![\text{likes}]\!]$.
- Given our syntax, we know that $[\![\textit{Bob}]\!]$ is going to combine with $[\![\textit{likes}]\!]$, so $[\![\textit{likes}]\!]$ must be a function which is applied to an argument to its right, viz. $[\![\textit{likes}]\!]([\![\textit{Bob}]\!])$
- We will therefore adopt the following composition rule:

Compared to (S1), notice that the combinatory order is reversed.

RULE S$_6$

If α has the form $\begin{array}{c}\text{VP}\\ \diagup \quad \diagdown\\ \beta \qquad \gamma\end{array}$ then $\alpha = [\![\beta]\!](\gamma)$

- Applying (S6), we get:

$$\left[\!\!\left[\begin{array}{c}\text{V}\\ |\\ \textit{likes}\end{array}\right]\!\!\right]\left(\left[\!\!\left[\begin{array}{c}\text{NP}\\ |\\ \text{N}\\ |\\ \textit{Bob}\end{array}\right]\!\!\right]\right)$$

- Applying (S2) and (S4) to the NP-node, and (S5) to the V-node, we get:

$$[\![\textit{likes}]\!]([\![\textit{Bob}]\!])$$

- It now follows:
  - $[\![\textit{likes}]\!]$ must be a function that which when applied to $[\![\textit{Bob}]\!]$ outputs a function that can then be applied to $[\![\textit{Mia}]\!]$.
  - This is only possible if $[\![\textit{likes}]\!]$ is a function from individuals to another function.

- I.e. $[\![likes]\!]$ must denote a function f which takes an individual, and then spits out a new function f' where f' is a function from individuals to truth values.

- And in conclusion:

  - $$[\![likes]\!] = \mathcal{I}(likes) = \Big[\ f\colon D_e \longmapsto \{g \mid g\colon D_e \longmapsto D_t\} \\ \text{for all } x \in D_e\colon f(x) = g_x\colon D_e \longmapsto D_t \\ \text{for all } y \in D_e,\ g_x(y) = 1 \text{ iff } y \text{ likes } x\ \Big]$$

- Hence, $[\![likes]\!]$ is a function from D_e to a function from D_e to D_t, viz. the characteristic function of $[\![likes]\!] \subseteq D_e{\times}(D_e{\times}D_t)$.

4.3.1 1-Place Function Valued Functions vs. 2-Place Functions

- Again, our semantics contrasts that of FOL. In FOL, 2-place predicates, e.g. 'like', denote sets of ordered pairs, viz. $[\![like]\!] \subseteq D_e{\times}D_e$.
- However, again, from a set of ordered pairs A, e.g. $\{\langle x,y\rangle \mid x \text{ likes } y\}$, where $A \subseteq D_e{\times}D_e$, we can determine a characteristic function. There is however one problem.
  - The characteristic function f^A of A (over $D_e{\times}D_e$) is going to be a subset of $(D_e{\times}D_e){\times}D_t$.
  - But if $f^A \subseteq (D_e{\times}D_e){\times}D_t$ then f^A is a function from *pairs of individuals* to truth values—which is not quite what we want.
  - What we want is a function from individuals to a new function from individuals to truth values. I.e. a function $f^* \subseteq D_e{\times}(D_e{\times}D_t)$.
  - Luckily, Moses Schönfinkel (and later Howard Curry) proved that 2-place functions can be reduced to 1-place functions. This process is normally referred to as **currying** (or **Schönfinkelization**).
- Let $D_e = \{\text{Mary, Bob}\}$ and suppose we have the following 2-place function:

$$\boldsymbol{f\colon D_e{\times}D_e \longmapsto D_t}$$

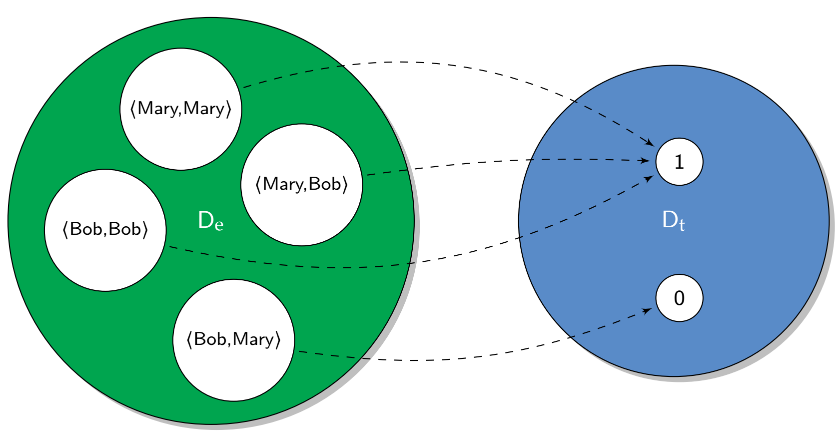

- This function denotes the following set of pairs.

$$\{\ \langle\langle \text{Mary,Mary}\rangle,1\rangle,\quad \langle\langle \text{Mary,Bob}\rangle,1\rangle,\quad \langle\langle \text{Bob,Mary}\rangle,0\rangle,\quad \langle\langle \text{Bob,Bob}\rangle,1\rangle\ \}$$

- We use right-to-left currying because the verb first combines syntactically with its direct object and then subsequently combines with the subject.

Right-to-Left Currying

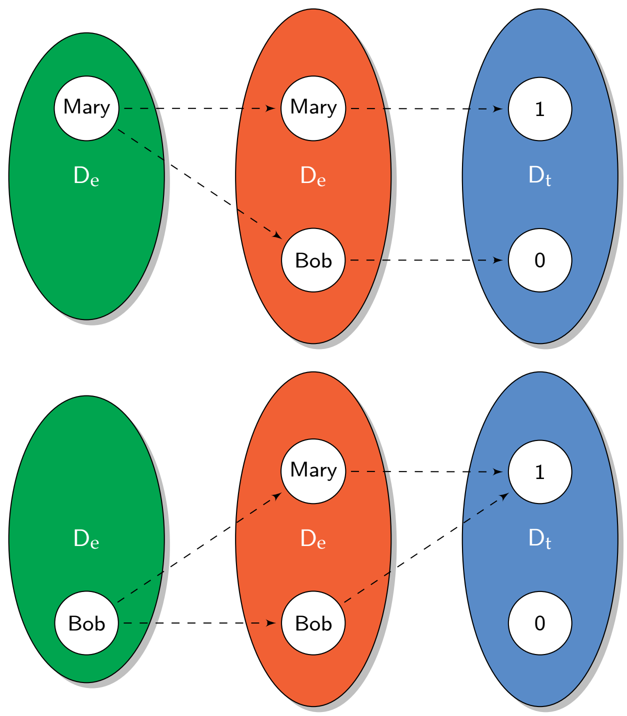

- This yields the following set of ordered pairs:

$$\{\ \langle \text{Mary},\langle \text{Mary},1\rangle\rangle,\quad \langle \text{Mary},\langle \text{Bob},0\rangle\rangle,\quad \langle \text{Bob},\langle \text{Bob},1\rangle\rangle,\quad \langle \text{Bob},\langle \text{Mary},1\rangle\rangle\ \}$$

- Here, the first coordinate of the ordered pair is object of the liking (the likee). The second coordinate is an ordered pair whose first coordinate is the individual doing the liking (the liker), and the second coordinate is the truth value.
- The point here is that any n-place function can be reduced to a 1-place (function valued) function.
- Transitive verbs, on our analysis, thus end up denoting functions $f \subseteq D_e \times (D_e \times D_t)$.
- But this is harmless, because there is a route from function-valued functions back to 2-place functions.

4.4 Type Theory

- The **theory of types** is a logical system, originally developed by Bertrand Russell to solve the set theoretic paradox known as Russell's Paradox.
- While type theory is no longer standard in set theory, type theory is a core tool in formal semantics.
- In type theory, one starts by characterizing a type theoretic language. One assumes that there is a set of types **T**. This set contains two basic types and these are then used to recursively define other complex types.
  - **Type *e***
    e is the type of individuals — so, D_e is a set of objects of type e.
  - **Type *t***
    t is the type of truth values — so, D_t is a set of objects of type t.
- The basic types correspond to the objects that Frege took to be *saturated*. From these basic types, complex types are recursively defined.
  1. $e \in \mathbf{T}$
  2. $t \in \mathbf{T}$
  3. If $\sigma \in \mathbf{T}$ and $\tau \in \mathbf{T}$, then $\langle \sigma, \tau \rangle \in \mathbf{T}$.
  4. Nothing is an element of **T** except on the basis of (i), (ii), and (iii).
- With this definition, we can now define an infinity of different types. For example:

$$\langle e,t \rangle \quad | \quad \langle e, \langle e,t \rangle\rangle \quad | \quad \langle t,t \rangle \quad | \quad \langle \langle e,t \rangle, t \rangle$$

We will soon discover that for each of these types there is an expression in English with that type.

- Complex types correspond to what Frege took to be *unsaturated* objects, i.e. *functions*. For example:
  - Expressions of type $\langle e,t \rangle$ are functions from objects of type e to objects of type t.
  - Expressions of type $\langle e, \langle e,t \rangle\rangle$ are functions from objects of type e to a function of type $\langle e,t \rangle$.
- With types as part of our inventory, we can now classify various expressions according to their type. For example.
  - $[\![Bob]\!]$ is type e
  - $[\![snore]\!]$ is type $\langle e,t \rangle$
  - $[\![like]\!]$ is type $\langle e, \langle e,t \rangle\rangle$
- The denotations of these expressions obviously depend on the model $\mathfrak{M}$, but they each have different domains — we will use the following convenient abbreviations for these domains.

DOMAINS	DESCRIPTION	SET THEORETIC OBJECTS
D_e	Set of individuals/objects x	$x \in D_e$
$D_{\langle e,t \rangle}$	Set of functions f from D_e to D_t	$f \subseteq D_e \times D_t$
$D_{\langle e, \langle e,t \rangle\rangle}$	The set of functions f from D_e to $D_{\langle e,t \rangle}$	$f \subseteq D_e \times (D_e \times D_t)$

- More generally, we will say that for any types σ and τ:

DOMAINS	DESCRIPTION	SET THEORETIC OBJECTS
$D_{\langle\sigma,\tau\rangle}$	Set of all functions from D_σ to D_τ	$f \subseteq \sigma \times \tau$

The Convenience of Types

- With type assignments to expressions of natural language, it is much easier to determine the semantics of new expressions.
- Suppose we want to figure out what semantics to give for the preposition 'to' — as used in the sentence (33).

(33) Sue talked to Bob.

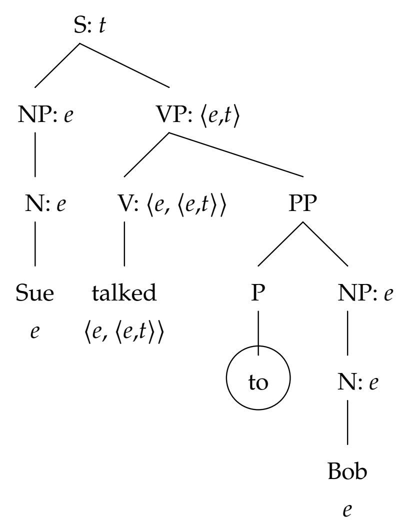

- Annotating the phrase structure with types makes it extremely easy to tell what type of expression is needed for the derivation to work out.

4.4.1 Simplifying Function Notation: λ-Expressions

- As functions get more involved, our notation will become more and more convoluted. We will remedy this problem by introducing λ-terms (lambda terms).
- λ-terms usually follow the schema below.

$$\lambda\alpha\colon \phi \;.\; \gamma$$

  - α is the **argument variable** (arbitrary letter standing for the argument of the function we are defining.)
  - ϕ is the **domain condition** (the domain over which the function is defined).

- γ is the **value description** (specification of the value (or output) of the function we are defining.)

- Here is an example of a λ-term.

$$\lambda x \colon x \in \mathbb{N} \,.\, x + 1$$

  - x is the **argument variable**.
  - $x \in \mathbb{N}$ is the **domain condition** (the set of natural numbers).
  - $x + 1$ is the **value description**.

- Since λ-terms are functions, these can be applied to arguments. Here is an example of the λ-term above as applied to an argument.

	$[\lambda x \colon x \in \mathbb{N} \,.\, x + 1](2)$
=	3

Notice how right-to-left currying becomes completely transparent when we use λ-terms.

- λ-terms makes it much less cumbersome to write down lexical entries. For example:
  - $[\![snore]\!] = [\lambda x \colon x \in D_e \,.\, x \text{ snores}]$
  - $[\![like]\!] = [\lambda x \colon x \in D_e \,.\, [\lambda y \colon y \in D_e \,.\, y \text{ likes } x]]$

- Generally, to make notation even simpler, I will omit the domain condition and simply indicate the domain of the function using subscripted variables, i.e.
  - $[\![snore]\!] = [\lambda x_e \,.\, x \text{ snores}]$
  - $[\![like]\!] = [\lambda x_e \,.\, [\lambda y_e \,.\, y \text{ likes } x]]$

Be careful when λ-terms are followed by arguments — brackets indicate which argument each function must be applied to.

- With λ-terms in our inventory, derivations become much easier to read and do.

	EXPRESSION	RULE APPLICATION
	$[\lambda x_e \,.\, [\lambda y_e \,.\, y \text{ likes } x]](\text{Bob})(\text{Mary})$	
=	$[\lambda y_e \,.\, y \text{ likes Bob}](\text{Mary})$	(Function Application)
=	1 iff Mary likes Bob	(Function Application)

- In tree form, the derivation would proceed as follows:

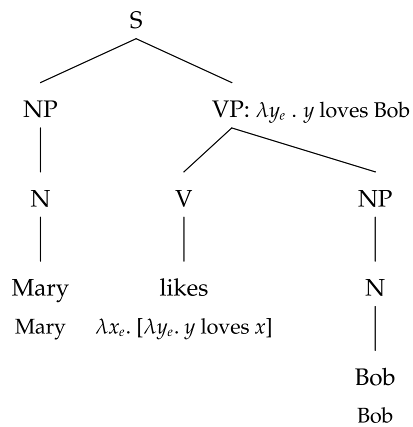

Important Note about λ-terms

- Without further formalization, our λ-terms are actually ambiguous — that is one can read $[\lambda \alpha_\phi \ . \ \gamma]$ in the following two ways:

  1. The function which maps every α such that ϕ to γ
  2. The function which maps every α such that ϕ to 1 if γ, and to 0 otherwise.

- For example, the denotation of an intransitive verb (cf. below) requires that the first λ-term is read as in 1. and the second λ-term is read as in 2.

$$\lambda x_e \ . \ [\lambda y_e \ . \ y \text{ loves } x]]$$

- The reason is that the first λ-term is a function from an individual to another function (a property). However, the second λ-term is a function from an individual to something that has a truth value.
- It will always be obvious which of the two readings is the correct one, so we will accept a bit of sloppiness and leave the λ-terms ambiguous.

4.5 Object Language vs. Meta-Language

- Another important distinction to keep in mind is that between object language and meta-language.

OBJECT LANGUAGE:
The object language is the language that is being analyzed — in our case English.

METALANGUAGE:
The metalanguage is the language that we use to study the object language. In our case, the metalanguage is richer than the object language in that it contains both English but also type theory, set theory and the λ-terms.

- Notice that the metalanguage assumed is quite powerful — significantly more powerful than e.g. FOL.

Distinguishing Object-Language from Meta-Language

- To indicate that an expression is part of the **object language**, one standardly encloses the expression in quotes.

'Barack Obama' 'Grass is green' 'snore'

- Also, remember the domain of the interpretation function $[\![\cdot]\!]$ is:

(a) Basic expressions of the object language.
(b) Complex expressions of the object language which satisfy both (i.) and (ii.) below
   i. They are the output of an accepted syntactic operation.
   ii. They are governed by already defined composition rules.

- This means that each of the examples below are undefined nonsense:

$\cancel{[\![\{1,2\}]\!]}$ $\cancel{[\![\text{adsfadsf}]\!]}$ $\cancel{[\![\text{grass green eat jumps}]\!]}$ $\cancel{[\![\lambda x\,.\,x \text{ snores}]\!]}$

- Correspondingly, when an expression is not enclosed in quotes (or otherwise indicated to be part of the object language), we are then **using** (rather than **mentioning**) the expression — i.e. we are using the expression its **denotation** and not the expression itself.
- Here's an example:

(a) Obama consists of five letters (b) 'Obama' consists of five letters.

- (a) is straightforwardly false. The president of the United States (i.e. Barack Obama) does not consist of five letters since he does not consist of letters at all.
- In contrast, (b) is true. The **expression** 'Obama' does consists of five letters.
- In these notes, I will enclose expressions in quotes when they are merely mentioned. If an expression is enclosed in $[\![\cdot]\!]$, I will italicize the expression to indicate that it is part of the object language.

Quiz Question: What is orange and rhymes with 'parrot'?

4.6 Type Driven Interpretation

- The composition rules (S1)–(S6) depend explicitly on syntactic categories, i.e. VP, V, NP, N, etc. But there are two principal reasons to be skeptical about this approach.

  a. The composition rules are highly dependent on the syntactic categories. Less syntactic-category driven rules will decrease the dependence on syntactic assumptions which are often in dispute.

  b. With this approach, we will need an implausibly large number of composition rules. Hence, for reasons of theoretical simplicity, more general rules would be better.

- To replace the composition rules above, we introduce the following **interpretation** rules.
- These rules are defined purely in terms of semantic types and only require that sentences are binary structured.

TERMINAL NODES (TN)
If α is a terminal node, $[\![\alpha]\!]$ is specified in the lexicon (i.e. if α is a basic expression of the object language, $[\![\cdot]\!]$ is defined for α).

NON-BRANCHING NODES
Non-Branching Nodes (NN)
If α is a non-branching node, and β is its daughter node, then $[\![\alpha]\!] = [\![\beta]\!]$.

FUNCTIONAL APPLICATION
If α is a branching node and $\{\beta,\gamma\}$ is the set of α's daughters, then if $[\![\beta]\!]$ is a function whose domain contains $[\![\gamma]\!]$, $[\![\alpha]\!] = [\![\beta]\!]([\![\gamma]\!])$.

- These rules capture everything that our previous rules did, but do not make any explicit reference to the syntactic categories associated with each constituent of the phrase structure tree. The rules only make reference to hierarchical structure — i.e. what dominates what.

Note the order insensitivity in the rule for functional application.

- Note that **Functional Application** now effectively has two meanings:

  (a) The process of applying functions to arguments — we now refer to this as β-reduction.

  (b) The interpretation rule that licenses the interpretation of a branching node as function and argument respectively — referred to as (FA).

Type Driven Interpretation: An Example

- Here is an example of a type driven interpretation of the sentence (34).

(34) Oswald killed Kennedy.

(35)

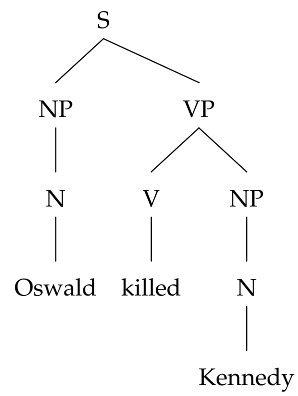

- We assume the following lexical entries:

  · $[\![\textit{Oswald}]\!]$ = Oswald

  · $[\![\textit{Kennedy}]\!]$ = Kennedy

  · $[\![\textit{killed}]\!] = [\lambda x_e\ .\ [\lambda y_e\ .\ y \text{ killed } x]]$

· Applying (FA) to (35).

$$[\![[_{VP} [_{V}\ \textit{killed}] [_{NP} [_{N}\ \textit{Kennedy}]]]]\!]\left([\![[_{NP} [_{N}\ \textit{Oswald}]]]\!]\right)$$

· Applying (NN) twice over to the NP-node.

$$[\![[_{VP} [_{V}\ \textit{killed}] [_{NP} [_{N}\ \textit{Kennedy}]]]]\!]([\![\textit{Oswald}]\!])$$

· Applying (FA) to the VP-node.

$$[\![[_{V}\ \textit{killed}]]\!]\left([\![[_{NP} [_{N}\ \textit{Kennedy}]]]\!]\right)([\![\textit{Oswald}]\!])$$

· Applying (NN) to the V-node and (NN) to the NP-node (twice over).

$$[\![\textit{killed}]\!]([\![\textit{Kennedy}]\!])([\![\textit{Oswald}]\!])$$

· From our lexicon and β-reductions.

$$\begin{aligned} & [\lambda x_e\ .\ [\lambda y_e\ .\ y \text{ killed } x]](\text{Kennedy})(\text{Oswald}) \\ = \; & [\lambda y_e\ .\ y \text{ killed Kennedy}](\text{Oswald}) \\ = \; & 1 \text{ iff Oswald killed Kennedy} \end{aligned}$$

Chapter 5

Non-Verbal Predicates and Modifiers

5.1 Vacuous Words

- There are many words in English whose contribution to truth conditions is somewhat unclear — for example the words 'is' and 'of' in (36) and the word 'a' in (37).

(36) Sam is proud of John.

(37) Whiskers is a cat.

(identity functions)

- To make things simpler, we will assume that these words make no essential contribution to the truth conditions. Given this, our analysis of these words should validate the following equalities.

(38) a. $[\![of\ John]\!] = [\![John]\!]$

b. $[\![is\ proud]\!] = [\![proud]\!]$

c. $[\![a\ cat]\!] = [\![cat]\!]$

- Suppose (36) has the phrase structure to the right and the following lexical entries.

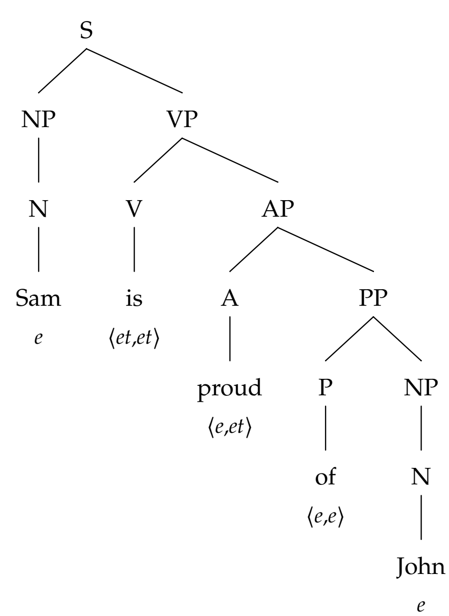

(39) $[\![proud]\!] = [\lambda x_e\ .\ [\lambda y_e\ .\ y \text{ is proud of } x]]$

(40) $[\![BE]\!] = [\lambda P_{\langle e,t\rangle}\ .\ P]$

(41) $[\![of]\!] = [\lambda x_e\ .\ x]$

(42) $[\![a]\!] = [\lambda P_{\langle e,t\rangle}\ .\ P]$

- An alternative to treating 'of' and 'is' as identity functions, is to simply assume that the syntax does not "see" these words.

5.1.1 Non-Verbal Predicates

- Simple predicates in the form of adjectives or nouns, e.g. 'is tired', 'is a cat', and 'is red', can now be treated just like intransitive verbs.

(43) $[\![tired]\!] = [\lambda x_e \, . \, x \text{ is tired}]$
(44) $[\![cat]\!] = [\lambda x_e \, . \, x \text{ is a cat}]$
(45) $[\![red]\!] = [\lambda x_e \, . \, x \text{ is red}]$

- Prepositions come in both intransitive and transitive varieties — but both can be represented along the lines of transitive and intransitive verbs.

(46) $[\![out]\!] = [\lambda x_e \, . \, x \text{ is not in } x\text{'s home}]$ (as in 'John is out')
(47) $[\![in]\!] = [\lambda x_e \, . \, [\lambda y_e \, . \, y \text{ is in } x]]$ (as in 'Sam is in Texas')

5.1.2 Restrictive Modifiers

- A Prepositional Phrase (PP) often occurs within an NP — for example.

(48) A city in Texas
(49) The barber from Seville
(50) Every sailor on the boat

- With regards to NP-modifying PPs, we have to distinguish between **restrictive** modifiers and **non-restrictive** modifiers.
- The examples in (48)–(50) are **restrictive** modifiers, because they function so as to restrict domain of the NPs that they combine with.
- (51) below is an example of a non-restrictive modifier (we won't concern ourselves with these here).

(51) Lance, (who is) from Texas, won the Tour de France.

The Problem with Restrictive Modifiers

- Consider the sentence below.

(52) Paris is a city in Texas.

Henceforth, I'll occasionally abbreviate the notation for types when convenient. For example, I will write $\langle et \rangle$ for functions of type $\langle e,t \rangle$, and I'll write $\langle et,et \rangle$ for functions of type $\langle \langle e,t \rangle,\langle e,t \rangle \rangle$.

(53)

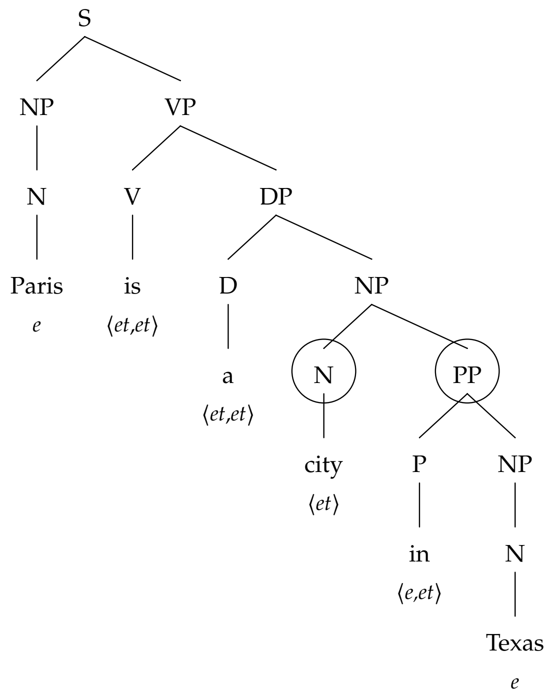

- If we try to interpret the PP- and N-node, we get a type-clash. This structure is therefore ruled out as uninterpretable. Since it should not be, we need e.g. another interpretation rule to deal with these cases.

> PREDICATE MODIFICATION (PM)
> If α is a branching node, and $\{\beta,\gamma\}$ is the set of α's daughters, and $[\![\beta]\!]$ and $[\![\gamma]\!]$ are both in $D_{\langle e,t\rangle}$, then: $[\![\alpha]\!] = [\lambda x_e \, . \, [\![\beta]\!](x) = 1 \text{ and } [\![\gamma]\!](x) = 1]$

- Applying PM to an example:

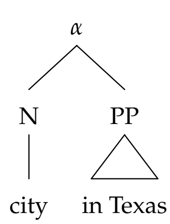

$$= [\lambda x_e \, . \, [\![\textit{city}]\!](x) = 1 \text{ and } [\![\textit{in Texas}]\!](x) = 1]$$

$$= [\lambda x_e \, . \, [\![\textit{city}]\!](x) = 1 \text{ and } [\lambda z_e \, . \, z \text{ is in Texas}](x) = 1]$$

$$= [\lambda x_e \, . \, [\lambda y_e \, . \, y \text{ is a city}](x) = 1 \text{ and } [\lambda z_e \, . \, z \text{ is in Texas}](x) = 1]$$

$$= [\lambda x_e \, . \, [\lambda y_e \, . \, y \text{ is a city}](x) = 1 \text{ and } x \text{ is in Texas}]$$

$$= [\lambda x_e \, . \, x \text{ is a city and } x \text{ is in Texas}]$$

- Could we do without PM?
- If $[\![city]\!]$ and $[\![in\ Texas]\!]$ is to combine via function application, the semantic types assigned to either 'city' or 'in Texas' must be revised.
- One option for dealing with NP-modification is thus the following:
  - $[\![in\ Texas]\!] = [\lambda P_{\langle et \rangle} \,.\, [\lambda y_e \,.\, P(y)$ and y is Texas]] $\langle et,et \rangle$
  - $[\![city]\!] = [\lambda x_e \,.\, x$ is a city] $\langle et \rangle$
- However, this now requires a different type assignment to the preposition 'in' —previously we assumed it was of type $\langle e,et \rangle$. One option:
  - $[\![in]\!] = [\lambda x_e \,.\, [\lambda P_{\langle et \rangle} \,.\, [\lambda y_e \,.\, P(y)$ and y is in x]]] $\langle e,\langle et,et \rangle\rangle$
  - $[\![Texas]\!]$ = Texas e
- For adjectives used to modify a noun, i.e. 'gray' in 'gray cat', we can do the same.
  - $[\![gray]\!] = [\lambda P_{\langle et \rangle} \,.\, [\lambda x_e.\ x$ is gray and x is P]] $\langle et,et \rangle$
  - $[\![cat]\!] = [\lambda y_e \,.\, y$ is a cat] $\langle et \rangle$
- But now we run into a problem with simple predicative uses of adjectival phrases and prepositional phrases, e.g. (54) and (55).

(54) Whiskers is gray.
(55) Paris is in Texas.

- **Further options**:
  a. Change the denotation for the copula 'be' (this adds other complications that we won't explore here).
  b. Assume that all adjectives are systematically ambiguous between e.g. $\langle et \rangle$-denotations and $\langle et,et \rangle$-denotations. Make similar assumption about prepositions.
  c. Assume that the default type of all adjectives is $\langle et,et \rangle$, and that these combine with a silent noun (a phonologically null element) when occurring in predicative position. (Montague's view).
- Neither of these options is clearly better than introducing PM, so we will stick with PM.

5.1.3 Predicate Modification and Non-Intersective Adjectives

- With our current rule of PM, we predict the following equivalence.

(56) Paris is a city in Texas ⇔ Paris is a city and Paris is in Texas

- This might seem OK for examples such as (55). However, it seems problematic for examples such as (57).

(57) Barbar is a large elephant.
$\not\Leftrightarrow$ Barbar is large and Barbar is an elephant.

- Many adjectives in English are **non-intersective**, i.e. the individuals that are large elephants are not co-extensive with the individuals that are in the intersection of the set of large things and the set of elephants.
- Non-intersective adjectives in English include *small*, *tall*, *fast*, *slow*, *smart*, *dumb* etc.

(examples of non-intersective adjectives)

- Non-intersective adjectives are **context sensitive** (or **context dependent**). Whether some individual x counts as *small*, *tall*, *fast*, *slow*, *smart*, *dumb* seems to depend on which comparison class is operative in the context.
- **Examples**

Context A
A group of biologists are trying to classify animals native to the savanna north of the Congo into three different groups, namely *large*, *intermediate*, and *small*. One biologist looks at Barbar and says "Barbar is small" — *false*.

Context B
A group of biologists are trying to classify elephants native to the savanna north of the Congo into three different groups, namely *large*, *intermediate*, and *small*. One biologist looks at Barbar and says "Barbar is small" — *true*.

5.1.4 More Difficult Cases

- We have covered both intersective and non-intersective adjectives. There is however a group of adjectives which pose even worse problems. Consider the sentences below.

(58) Bob is an alleged criminal.
(59) Sue is a former republican.
(60) Sam bought a fake gun.

- These adjectives appear to explicitly preclude membership in the set of things denoted by the noun with which they combine, i.e.

— a is an alleged criminal $\not\models$ a is a criminal.

— a is a former republican $\not\models$ a is a republican.

— a owns a fake gun $\not\models$ a owns a gun.

- In short, the meaning of e.g. the NP 'a fake gun' should be explicated in such a way that its denotation is not a member of the set of guns.
- How to deal with these cases would not only require a revision of our semantics, it would require answering some deep metaphysical questions about what it takes for something to be e.g. a fake gun.

Chapter 6

The Definite Article

6.1 Definite Descriptions

- Intuitively, expressions such as (61) and (62) (called definite descriptions) simply refer to whoever the descriptions pick out.

  (61) The president of the United States
  (62) The lecturer from Denmark

- A definite description is composed of a determiner (D) and a noun phrase (NP).

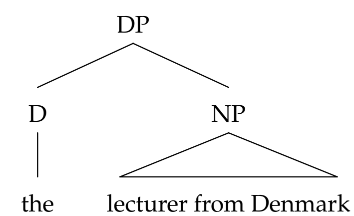

- Since an expression of the form 'lecturer from Denmark' is a property (a 1-place predicate), its denotation should be the following.

$$[\![\textit{lecturer from Denmark}]\!] = \quad \lambda x_e \,.\, x \text{ is a lecturer from Denmark} \qquad \langle et \rangle$$

- Given this, one would think that $[\![\textit{the}]\!]$ should be a function which takes an expression of type $\langle et \rangle$ (a property) as input and outputs an individual e (the individual who has that property) — viz. a function of type $\langle et{,}e \rangle$

6.1.1 Russell's Analysis

- Russell (1905) observed that there are obvious syntactic similarities between a number of different expressions in English, e.g.

(63) The Russian spy.
(64) A/Some Russian spy.
(65) Every Russian spy.
(66) Most Russian spies.
(67) Both Russian spies.
(68) Three Russian spies.

- I.e. these expressions appear to serve the same syntactic purpose.

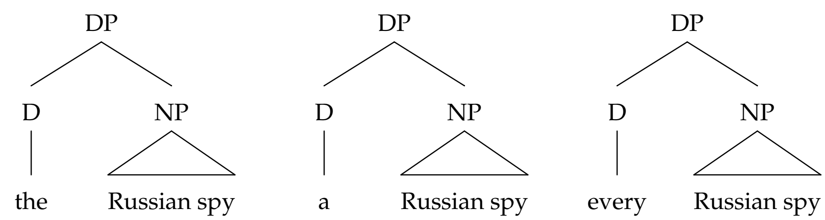

- Russell argued that these expressions form one *class* of expressions (which he called *denoting phrases*). Today this group of expressions are referred to as *determiner phrases.*

Syntactic and Semantic Uniformity

- Given the syntactic uniformity of these expressions, one might think that our semantic analysis should also, to some extent, be uniform.
- However, as Russell observed, determiner phrases cannot generally be analyzed as referential terms (expressions of type e). For example, an expression of the form 'a Russian spy' cannot be treated as referring to a specific individual.

(69) A Russian spy killed Kennedy.

- The sentence in (69) intuitively is true as long as some Russian spy or other killed Kennedy (not only if some *specific* Russian spy killed Kennedy).
- Russell therefore argued that an indefinite description such as 'a spy from Russia' should be analyzed as a type of quantifier phrase, i.e.

(69) A Russian spy killed Kennedy.
   a. There is an individual x: x is a Russian spy and x killed Kennedy.
      $\exists x[\text{Russian-spy}(x) \wedge \text{killed}(x, \text{Kennedy})]$

- If determiners in general should be given a uniform treatment, analyzing definite descriptions as some sort of quantifier phrase seems natural.

Uniqueness

- There is a clearly a difference in meaning between definite and indefinite descriptions, so in what way should their semantic analysis differ?

(69) A Russian spy killed Kennedy.
(70) The Russian spy killed Kennedy.

- (70) seems to imply (in some sense) that there is a *unique* Russian spy — (69) does not. Russell therefore proposed to analyze (70) as follows.

(70) The spy from Russia killed Kennedy.

a. There is exactly one Russian spy and that spy killed Kennedy.
$\exists x[\text{Russian-spy}(x) \wedge \forall y[\text{Russian-spy}(y) \rightarrow y = x \wedge \text{killed}(x, \text{Kennedy})]]$

- Notice, crucially, that on this analysis definite descriptions do not **refer**. That is, definite descriptions are not **referential terms**.

Further Motivation for Russell's Analysis: Non-Denoting Descriptions

- Suppose that definite descriptions are referential terms. If so, what is the referent of the following definite descriptions.

(71) The king of France shot my cat last night.
(72) The planet between the earth and the moon is inhabitable.
(73) Kate Middleton is married to the Duke of Cambridge or the Emperor of Oxford.

- Since the definite descriptions in these sentences do not refer to anything, it is not clear that the sentences have truth conditions. However, if the sentences have no truth conditions, then we effectively predict that they are meaningless. That doesn't seem right.
- Russell's analysis provides a solution to this problem: The meaning of (71) is paraphrased below.

(74) There is exactly one x such that x is a king of France and x shot my cat last night.
I.e. $\exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y \wedge \text{shot-my-cat-last-night}(x)]]$

6.1.2 A Russellian Semantics for Definite Descriptions

- In our current semantic system, we can implement the Russellian analysis as follows.

(75) $[\![\textit{the}]\!]^R = [\lambda P_{\langle et \rangle} \,.\, [\lambda Q_{\langle et \rangle} \,.\, \text{there is exactly one } x \text{ who is } P \text{ and } x \text{ is } Q]]$ $\langle et, \langle et,t \rangle\rangle$

- Equivalently, but a bit more elegant.

(76) $[\![\textit{the}]\!]^R = [\lambda P_{\langle et \rangle} \,.\, [\lambda Q_{\langle et \rangle} \,.\, \exists x\big(\mathrm{P}(x) \wedge \forall y\big(P(y) \rightarrow x = y \wedge Q(x)\big)\big)]]$ $\langle et, \langle et,t \rangle\rangle$

- To see that this would work consider the following.

(77)

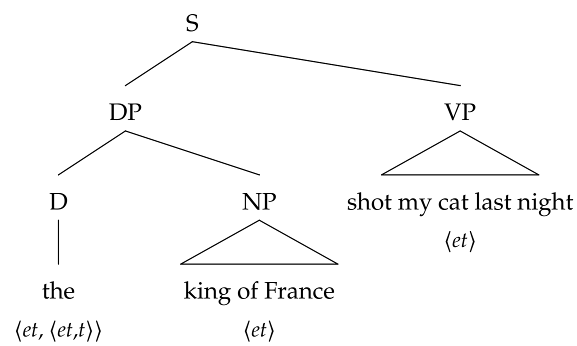

- On this analysis of 'the *F*', the determiner phrase is being treated as **generalized quantifier** — and we will talk more about such expressions in the coming weeks.
- For now, we will consider an alternative to Russell's analysis.

6.1.3 The Frege-Strawson Analysis

- A famous quote from Strawson (1950).

> Now suppose some one were in fact to say to you with a perfectly serious air: "The king of France is wise". Would you say, "That's untrue"? I think it's quite certain that you wouldn't. But suppose he went on to ask you whether you thought that what he had just said was true, or was false; whether you agreed or disagreed with what he had just said. I think you would be inclined, with some hesitation, to say that you didn't do either; that the question of whether his statement was true of false simply *didn't arise*, because there was no such person as the king of France. You might, if he were obviously serious (had a dazed astray-in- the-centuries look), say something like: "I'm afraid you must be under a misapprehension. France is not a monarchy. There is no king of France." (1950, 330)

- Here Strawson is arguing that if a speaker uses a description that fails to refer, the speaker appears to be making some kind of **linguistic** mistake rather than saying something false (as Russell's analysis would predict).
- Before Strawson, Frege (1892) argued for a similar view for both non-referring names and descriptions. Here's Frege on proper names.

> If therefore one asserts 'Kepler died in misery', there is a presupposition that the name 'Kepler' designates something; but it does not follow that the sentence 'Kepler died in misery' contains the thought that the name 'Kepler' designates something. If this were the case the negation would have to run not
>
> 'Kepler did not die in misery',
>
> but
>
> 'Kepler did not die in misery, or the name "Kepler" is *bedeutungslos*'.
>
> (Frege, 1892, 163-163)

- According to the Frege–Strawson view...

  (a) 'the F is G' does not assert, but rather **presupposes**, the existence of a unique F.

  (b) 'the F is G' is not truth evaluable if there is no unique F.

- **Other arguments against Russell's analysis**
  Russell argued that the meaning of (78) is actually (79) (roughly speaking).

  (78) The king of France is bald.
  (79) $\exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y \wedge \text{bald}(x)]]$

- If Russell is right, it seems hard to explain why a speaker who utters (80)–(81) cannot thereby mean (82) — and similarly for (83) and (84).

  (80) Probably, the king of France is bald
  (81) It is probable that the king of France is bald
  (82) Probably, $\exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y \wedge \text{bald}(x)]]$.

  (83) If the king of France is bald, then Obama is bald too.
  (84) $\exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y \wedge \text{bald}(x)]] \rightarrow \text{bald}(\textit{Obama})$

A Presuppositional Analysis of Definite Descriptions

- **Partial Functions**
  How can we capture formally this idea that definite descriptions trigger a presupposition? One option is to define 'the' as a partial function.

  Recall, a function $f\colon A \longmapsto B$, is *partial* if it maps only *some* of the elements of A to B but never maps any element of A to more than one element in B.

  (85) $[\![the]\!]^{\text{FS}} = [\lambda P_{\langle et \rangle}$: **there is exactly one x such that $P(x)$** . the unique x such that $P(x)]$

- The description in boldface (between the colon and the period) is meant to indicate that this function is partial. It is defined only for properties P that satisfy the requirement that P has exactly one member.

- More conspicuously...

  (86) $[\![the]\!]^{\text{FS}} = [\lambda P_{\langle et \rangle}\colon \exists! x P(x) \,.\, \iota x P(x)]$

  A function from properties P—defined only if there is exactly one P—to the unique individual x who is a member of P.

- The ι-term, $\iota x P(x)$ is a name in the metalanguage for the unique individual who is P. So, it is a referential term functioning essentially like an individual constant in predicate logic — it simply refers to the unique individual who is P.

- Putting the Frege-Strawson analysis to work...

(87)

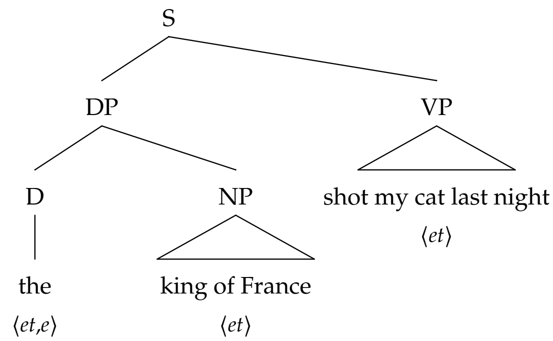

6.1.4 Presuppositions as Partial Functions: Some Consequences

- While there are several compelling reasons in favor of the Frege–Strawson analysis, the analysis comes with a price.
- For example, if sentences with non-denoting descriptions are predicted to have no truth value (i.e. be neither true nor false), this has serious ramifications for our semantic system, namely that the system is no longer bivalent. That is, not every sentence has either the value *true* or the value *false*.
- In a non-bivalent semantic system, the meanings of various sentential (logical) connectives, e.g. conjunction, disjunction, conditionals, cannot be defined as usual.
- Now, according to the Frege–Strawson analysis, (88) is neither true nor false.

(88) The king of France is bald.

- So, what would now be the correct prediction for sentences such as (89)–(92)?

(89)	2+2=5 and the king of France is bald.	$(F \wedge ?)$
(90)	2+2=4 and the king of France is bald.	$(T \wedge ?)$
(91)	The king of France is bald or the queen of Denmark is a liar.	$(? \vee F)$
(92)	If the king of France is bald, the queen of Denmark is not a liar.	$(? \rightarrow T)$

- Should these also be predicted to be neither true nor false? Maybe so.

Notice—for example—that in classical logic (93) is an instance of the law of excluded middle, ($P \vee \neg P$), a logical truth.

- But what about these (93)–(97).

(93)	The king of France is bald or the king of France is not bald.	$(? \vee ?)$
(94)	If there is a king of France, the king of France is bald.	$(F \rightarrow ?)$
(95)	If there were a king of France, the king of France would be bald.	$(F \mathrel{\Box\!\!\rightarrow} ?)$
(96)	There is a king of France and the king of France is bald.	$(F \wedge ?)$
(97)	Either there is no king of France or the king of France is in hiding.	$(T \vee ?)$

6.1.5 A Few Remarks on Uniqueness

Remember, on Russell's analysis, uniqueness is what distinguishes indefinites from definites!

- Both Russell's analysis and the Frege–Strawson analysis makes a uniqueness assumption: 'the *F*' refers/denotes a unique individual who is *F*.
- This particular assumption is not unproblematic, because there are numerous situations in which it seems that an utterance of 'the *F* is *G*' would be true, but where the uniqueness condition is not satisfied.
- Consider the following:

(98) The table is covered with books.
  a. The table [**in this room**] is covered with books.

(99) The president passed the bill.
  a. The [**US**] president passed the *health care* bill.

(100) Unsurprisingly, the Russian voted for the Russian.
  a. Unsurprisingly, the Russian [**referee**] voted for the Russian [**boxer**].

- In other words, it seems that definite descriptions are often used to convey more specific information than that intuitively contained in the sentence itself.
- However, it is important to realize that this is not a problem exclusive to definite descriptions.

(101) Every table is covered with books.
  a. Every table [**in this room**] is covered with books.

(102) Unsurprisingly, most Russians voted for the Russian.
  a. Unsurprisingly, most Russians [**voters**] voted for the Russian [**candidate**].

(103) Some students passed less than one exam.
  a. Some students [**at the University of Edinburgh**] passed less than one exam [**at the University of Edinburgh**]

- In other words, it seems that determiner phrases in general are sensitive to context (similar to the non-intersective adjectives discussed earlier).

Chapter 7

Relative Clauses and Variables

7.1 Relative Clauses

- Intuitively, the sentences in (104) and (105) have the same meaning. As a result, it seems reasonable to suppose that their meanings should be derived in similar ways.

(104) The empty house is available.
(105) The house which is empty is available.

- In (104), the semantic value of 'house' and 'empty' is computed using the rule of predicate modification, so we assume that in (105), 'house' and the relative clause 'which is empty' also combine via predicate modification. In short, our denotation for the relative clause 'which is empty' should end up denoting a property of type $\langle et \rangle$ — namely (106).

(106) $[\![\textit{which is empty}]\!] = [\lambda x_e \, . \, x \text{ is empty}]$

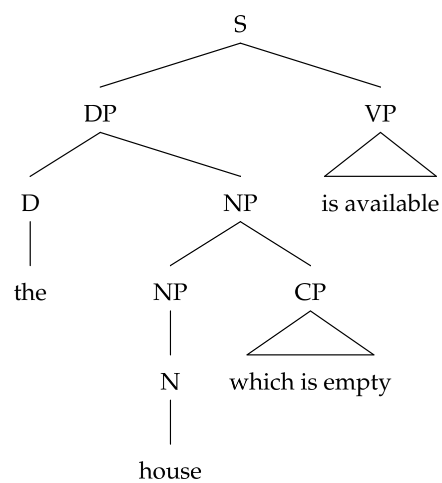

- The structure of (105) is illustrated by the phrase structure on the right.

- If we have the following lexical entries...

  $[\![house]\!] = [\lambda x_e \,.\, x \text{ is a house}]$

  $[\![which\ is\ empty]\!] = [\lambda x_e \,.\, x \text{ is empty}]$

- ... our derivation works out perfectly since 'house' and 'which is empty' can combine via PM.

7.1.1 Semantic Composition Inside the Relative Clause

- Our task now is to provide a systematic analysis of relative clauses which compositionally turn these into properties (i.e. expressions of type $\langle et \rangle$.)
- We need to derive the equivalence below.

$$[\![empty]\!] = [\lambda x_e \,.\, x \text{ is empty}] = [\![which\ is\ empty]\!]$$

Since we can have multiple embeddings of relative clauses, we enumerate the *wh*-expression and its corresponding trace to keep track.

- We assume that the internal structure of a restrictive relative clause looks as follows.

(107)

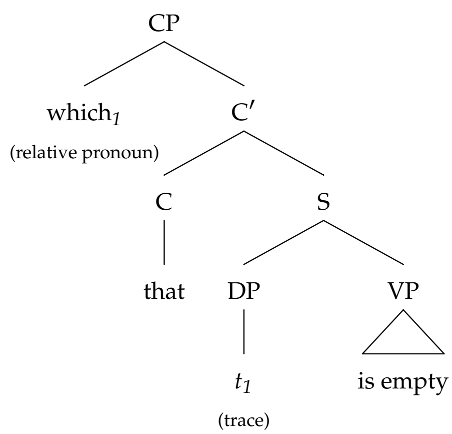

- A couple of notes about (107):
  - — Proper names, pronouns, traces are now assumed to be of syntactic category DP.
  - — Either 'which' or the complementizer *C* has to be deleted on the surface.

Passing VP-meaning up the tree.

- Since $[\![empty]\!]$ is a function of type $\langle et \rangle$ and its meaning is precisely what we want to end up with for (107), one might be tempted to simply assume that the meaning of (107) is just the meaning of the VP.
- However, this won't work in other cases, e.g.

(108) The house **which John abandoned**

(109)

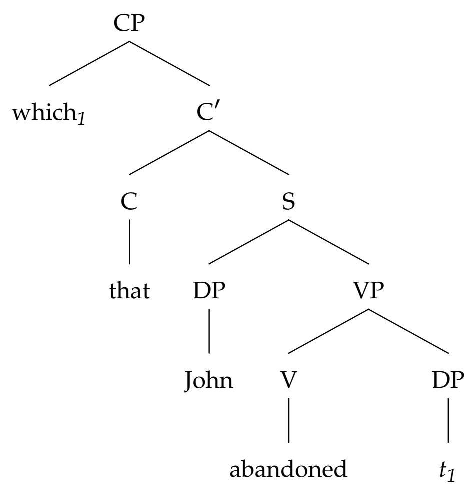

- Traces are variables so if we simply pass up the meaning of the VP here, we would end up with an ill-formed property (a property with an unbound variable in it).

What Do Traces Denote?

- Given that our structures contain traces t, we need to say something about the denotation of traces.
- However, given what the structures in (107) and (109) should ideally denote, the following *illustration* might provide a plausible suggestion.

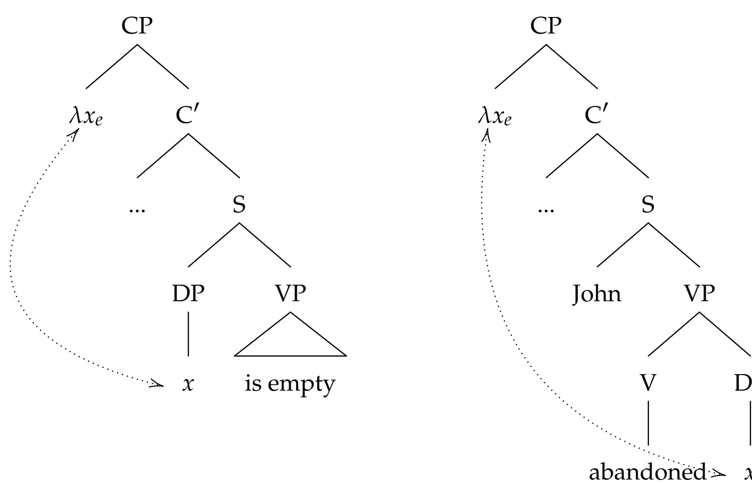

- If traces were treated as **variables** and these variables could, somehow, be bound by a λ-operator, that would give us exactly what we need.

- However, there is one obvious problem here. The above structures contain a λ-operator, but that is nonsensical: There are no λ-operators in English!
- In order to show how to get around this problem, we need to talk a bit about variables.

7.1.2 Variables and their interpretation

- In our recap of FOL, we saw that the extension of an expression is always relative to a model. E.g.

$$[\![\textit{Sue}]\!]^{\mathfrak{M}} = \text{Sue} \qquad [\![\textit{snore}]\!]^{\mathfrak{M}} = \{x \mid x \text{ snores}\}$$

- **Variables** (and hence **traces**) are essentially placeholders for individuals. However, these expressions do not receive their interpretation relative to the model (because what a variables denotes is not fixed once and for all by the model).
- One type of expression in natural language which both looks and feels like an unbound variable is a deictically used (or demonstratively used) pronoun, e.g.

(110) **She** snores. (111) **He** put **it** in the closet.

- These uses of pronouns contrast uses where the pronouns are *bound*, e.g.

(112) [Every philosopher]$_{1}$ thinks that he$_{1/2}$ is the smartest.
(113) [Every philosopher]$_{1}$ loves himself$_{1/*2}$.
(114) [Every philosopher]$_{1}$ loves him$_{*1/2}$.

- In FOL, sentences with an unbound variable (the natural language equivalent of a deictic pronoun) is interpreted using a **variable assignment**.
- With respect to our current formal system, we treat variable assignments g as partial functions from the natural numbers to individuals in D_e.

Why is the function partial? Because we want provisions for cases where no individual is picked out by the pronoun/trace, e.g. when there is no demonstration.

$$g \colon \mathbb{N} \longmapsto D_e$$

- For example, a variable assignment g might yield the following mapping from numbers to individuals.

$$g = \begin{bmatrix} 1 & \longmapsto & \text{Sue} \\ 2 & \longmapsto & \text{Bob} \\ 3 & \longmapsto & \text{Frank} \\ 4 & \longmapsto & \text{Mary} \\ 5 & \longmapsto & \text{Sue} \\ \vdots & & \vdots \end{bmatrix}$$

- We will generally assume that the context of utterance always determines a unique variable assignment. This means that our interpretation function should be relativized, not only to a model, but also to a variable assignment — i.e.

$$\cancel{[\![\alpha]\!]} \qquad\qquad [\![\alpha]\!]^{g}$$

- We adopt the following rule for the interpretation of variables (e.g. pronouns and traces).

(115) **Pronouns and Traces Rule**
If α is a pronoun or a trace, g is a variable assignment, and i is an index in the domain of g, **dom**(g), then $[\![\alpha_i]\!]^g = g(i)$

7.1.3 Assignment-Sensitive Interpretation Rules

- Since we are now assuming that some phrase structures will contain elements that only have extensions relative to variable assignments, we need to reformulate our interpretations rules.

- First we introduce the following definition of assignment-independent denotations.

(116) **Assignment Independent Denotations (AID)**
For any phrase structure α, α is in the domain of $[\![\cdot]\!]$ iff for all assignments g and g^*, $[\![\alpha]\!]^g = [\![\alpha]\!]^{g^*}$.

If α is in the domain of $[\![\cdot]\!]$, then for all assignments g, $[\![\alpha]\!] = [\![\alpha]\!]^g$.

That is, α is in the domain of the interpretation function (relative to model $\mathfrak{M}$, but NOT relative to a variable assignment g) if and only if the denotation of α is invariant to different assignments, e.g. g' and g''.

- Next, we reformulate our interpretation rules.

Terminals (TN)
If α is a terminal node occupied by a lexical item, then $[\![\alpha]\!]$ is specified in the lexicon.

Non-Branching Nodes (NN)
If α is a non-branching node and β its daughter, then, for any assignment g, $[\![\alpha]\!]^g = [\![\beta]\!]^g$.

Functional Application (FA)
If α is a branching node and $\{\beta, \gamma\}$ the set of its daughters, then, for any assignment g, if $[\![\beta]\!]^g$ is a function whose domain contains $[\![\gamma]\!]^g$, then $[\![\alpha]\!]^g = [\![\beta]\!]^g([\![\gamma]\!]^g)$.

Predicate Modification (PM)
If α is a branching node and $\{\beta, \gamma\}$ the set of its daughters, then, for any assignment g, if $[\![\beta]\!]^g$ and $[\![\gamma]\!]^g$ are both functions of type $\langle e,t\rangle$, then $[\![\alpha]\!]^g = [\lambda x_e \,.\, [\![\beta]\!]^g(x) = [\![\gamma]\!]^g(x) = 1]$.

- A sample derivation — suppose:

$$g = \begin{bmatrix} 1 \longmapsto \text{Sue} \\ 2 \longmapsto \text{Bob} \\ \vdots \longmapsto \vdots \end{bmatrix}$$

$[\![she_1\ left]\!]^g$	$= [\![left]\!]^g([\![she_1]\!]^g)$	(by FA)
	$= [\![left]\!]([\![she_1]\!]^g)$	(By AID)
	$= [\lambda x_e\ .\ x \text{ left}]([\![she_1]\!]^g)$	(by TN)
	$= [\lambda x_e\ .\ x \text{ left}](g(1))$	(By Pronouns and Traces Rule)
	$= [\lambda x_e\ .\ x \text{ left}](\text{Sue})$	
	$= 1$ iff Sue left	

7.1.4 Back to Restrictive Relative Clauses

- In the phrase structure below, we assume that the complementizer 'that' is semantically vacuous (we essentially act as if the syntax does not "see" this item).

(117)

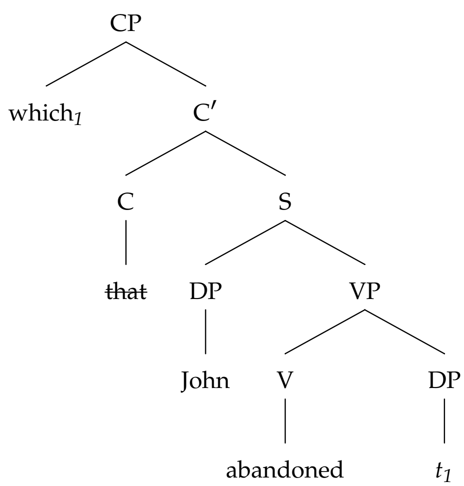

- The relative pronoun 'which' is assumed to have no denotation on its own, but we won't treat it as vacuous (as we did with the complementizer).
- Instead, we will introduce a rule which allows us to interpret these structures (and which makes direct reference to the relative pronoun).

(this is also called λ-abstraction)

> **Predicate Abstraction (PA)**
> If α is a branching node whose daughters are a relative pronoun with index i and β (where β contains a variable with index i), then $[\![\alpha]\!] = [\lambda x_e\ .\ [\![\beta_i]\!]^{g^{[x/i]}}]$

- $g^{[x/i]}$ is a **modified variable assignment**. It means that g has been modified so as to assign x to every occurrence of i in g.
- Predicate Abstraction (or λ-abstraction) is basically a rule that allows us form predicates from sentences with unbound variables (under certain specific circumstances).

- Here's a simpler way of thinking about what kind of operation predicate abstraction is.
  - Consider the language of FOL. In FOL, we have variables, predicates, quantifiers etc. Now, in FOL, our syntax permits the formation of expressions such as (118).

    (118) $F(x)$

  - However, the expression in (118) is only interpretable relative to an assignment of a value to the unbound variable. But, (118) could be made interpretable by performing a syntactic operation, namely binding the variable as in e.g. (119).

    (119) $\exists x F(x)$

  - Predicate abstraction is a similar type of syntactic operation. In particular, it is an operation that permits us to turn open sentences into properties by *abstracting* over the free variable (hence the terms *predicate abstraction* or λ-*abstraction*).

    (120) $\lambda x \,.\, F(x)$ (the characteristic function of F over D_e)

  - In a type theoretical language (which we are assuming is part of our metalanguage), variables are of type e, truth values are of type t.
  - So, strictly speaking, λ-abstraction is a syntactic rule that we can introduce into our type theoretical language $\mathcal{L}$, namely the following.

    (121) If α is an expression of type a in $\mathcal{L}$, and v is a variable of type b, then $\lambda v \alpha$ is an expression of type $\langle b,a \rangle$ in $\mathcal{L}$.

  - In the context of our current semantic system, the function of the predicate abstraction rule is to (a) introduce a λ-term and (b) *shift* the variable-assignment from assigning a particular individual to the trace t_i and instead replacing t_i with a variable which gets *bound* by the λ-term.

$$
\begin{aligned}
& \lambda y_e \,.\, [\![snores(t_i)]\!]^{g^{[y/i]}} \\
= \; & \lambda y_e \,.\, \text{snores}(g^{[y/i]}(t_i)) \\
= \; & \lambda y_e \,.\, \text{snores}(y)
\end{aligned}
$$

  - It is crucial to understand that the rule can only be used when the sentence in question has a very specific syntactic structure.
  - Remember λ-terms are not part of the object language. So, this is a syntactic operation in the metalanguage (which then has an effect on the object language).

7.2 A Derivation Using Predicate Abstraction

- Here is a semantic derivation of the relative clause in (108).

$$[\![\ [_{CP}\ \text{which}_1\ [_{C'}\ [_{C}\ \cancel{\text{that}}]\ [_{S}\ [_{DP}\ \text{John}]\ [_{VP}\ [_{V}\ \text{abandoned}]\ [_{DP}\ t_1]]]]]\ ]\!]^{g}$$

· Step 1: By Predicate Abstraction (PA).

$$[\lambda x_e\ .\ [\![\ [_{C'}\ [_{C}\ \cancel{\text{that}}]\ [_{S}\ [_{DP}\ \text{John}]\ [_{VP}\ [_{V}\ \text{abandoned}]\ [_{DP}\ t_1]]]]\ ]\!]^{g^{[x/1]}}]$$

· Step 2: By Vacuity of *C*.

$$[\lambda x_e\ .\ [\![\ [_{S}\ [_{DP}\ \text{John}]\ [_{VP}\ [_{V}\ \text{abandoned}]\ [_{DP}\ t_1]]]\ ]\!]^{g^{[x/1]}}]$$

· Step 3: By Functional Application (FA).

$$[\lambda x_e\ .\ \left[\!\!\left[\begin{array}{c}\text{VP}\\ \begin{array}{cc}\text{V} & \text{DP}\\ | & |\\ \text{abandoned} & t_1\end{array}\end{array}\right]\!\!\right]^{g^{[x/1]}}\ ([\![\mathit{John}]\!]^{g^{[x/1]}})]$$

· Step 4: By Assignment Independent Denotation (AID)

$$[\lambda x_e\ .\ \left[\!\!\left[\begin{array}{c}\text{VP}\\ \begin{array}{cc}\text{V} & \text{DP}\\ | & |\\ \text{abandoned} & t_1\end{array}\end{array}\right]\!\!\right]^{g^{[x/1]}}\ ([\![\mathit{John}]\!])]$$

· Step 5: By Lexical Entry for 'John'

$$[\lambda x_e\ .\ \left[\!\!\left[\begin{array}{c}\text{VP}\\ \begin{array}{cc}\text{V} & \text{DP}\\ | & |\\ \text{abandoned} & t_1\end{array}\end{array}\right]\!\!\right]^{g^{[x/1]}}\ (\text{John})]$$

· Step 6: By Functional Application (FA):

$$[\lambda x_e\ .\ \left[\!\!\left[\begin{array}{c}\text{V}\\ |\\ \text{abandoned}\end{array}\right]\!\!\right]^{g^{[x/1]}}\left(\left[\!\!\left[\begin{array}{c}\text{DP}\\ |\\ t_1\end{array}\right]\!\!\right]^{g^{[x/1]}}\right)(\text{John})]$$

- Step 7: By Non-Branching Nodes (NN) (twice over) — $[\lambda x_e \,.\, [\![abandoned]\!]^{g^{[x/1]}}([\![t_1]\!]^{g^{[x/1]}})(\text{John})]$

- Step 8: By Pronouns and Traces Rule (two steps) — $[\lambda x_e \,.\, [\![abandoned]\!]^{g^{[x/1]}}(g^{[x/1]}(t_1))(\text{John})]$
  $[\lambda x_e \,.\, [\![abandoned]\!]^{g^{[x/1]}}(x)(\text{John})]$

- Step 9: By Assignment Independent Denotations (AID) — $[\lambda x_e \,.\, [\![abandoned]\!](x)(\text{John})]$

- Step 10: By lexical entry for 'abandoned' — $[\lambda x_e \,.\, [\lambda y_e \,.\, [\lambda z_e \,.\, z \text{ abandoned } y]](x)(\text{John})]$

- Step 11: By β-conversion — $[\lambda x_e \,.\, [\lambda z_e \,.\, z \text{ abandoned } x](\text{John})]$

- Step 12: By β-conversion — $[\lambda x_e \,.\, \text{John abandoned } x]$

7.2.1 Defining Modified Variable Assignments

- Our rule of PA made use of a so-called modified variable assignment—indicated by the notation $g^{[x/i]}$.
- However, we have not formally defined this notion. We do that as follows:
  - Let g be a variable assignment, let $i \in \mathbb{N}$ and $x \in D_e$. If so, $g^{[x/i]}$ is the unique assignment which satisfies 1-3 below:
    1. $\mathbf{dom}(g) \cup \{i\} = \mathbf{dom}(g^{[x/i]})$
    2. $g^{[x/i]}(i) = x$
    3. For every $j \in \mathbf{dom}(g^{[x/i]})$ where $j \neq i$: $g^{[x/i]}(j) = g(j)$

Chapter 8

Quantifiers

8.1 Expressions of Generality

- Very roughly speaking, quantificational expressions are expressions of generality. Quantifiers are used to talk about quantities of things or amounts of stuff.
- In many languages, e.g. English, quantificational meanings are expressed using a determiner + noun. However, there are other languages, where quantificational meanings are expressed through other means, e.g. by agreement markers on verbs (ASL), by predicate modifiers or complementizers (Straits Salish).
- In these lectures, we will focus on:
  - Quantificational determiners in English that combine syntactically with *count* nouns (rather than *mass* nouns).
  - How to provide an adequate semantic analysis of expressions such as 'no *F*s', 'every *F*', and 'most *F*' etc.
- When talking about quantification in English, we need to distinguish between quantificational determiners and quantificational noun phrases.
  1. **Quantificational Determiners**: *some*, *every*, *most*, *many*, *three* etc.
     Quantificational determiners are of syntactic category D.
  2. **Quantified Noun Phrases (QNPs)**: *some boys*, *every girl*, *many tired professors* etc.
     We will assume that QNPs are of syntactic category DP and that they are compounds of a determiner D and a noun phrase NP.
- QNPs are expressions of *generality* — i.e. used to make general statements (and express general propositions). Some examples.

· Every bird flies	· Many birds fly
· Some bird flies	· Exactly five birds fly
· Two/six/seventeen birds fly	· More than two birds fly
· Both birds fly	· Several birds fly

Although, even in English, one can use some adjectives as quantifiers, e.g. 'numerous' and 'innumerably', and non-phrasal constructions like 'there is/are' also appear to be a type of quantification over objects. Adverbs like 'always', 'mostly', 'usually' are also standardly assumed to be quantifiers, albeit quantifiers over different types of entities.

- Integrating these expressions into our compositional semantics, we need to determine:
  (a) The denotation of QNPs (+ their semantic type).
  (b) The denotation of quantificational determiners (+ their semantic type).

8.2 The Denotation of Quantified Noun Phrases

8.2.1 A Problem with QNPs as *e*-Type expressions

- Since QNPs seemingly occur in the same syntactic positions as expressions of type e, cf. below, one might be tempted to treat them as type e.

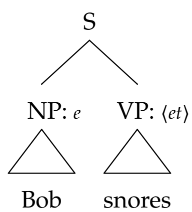

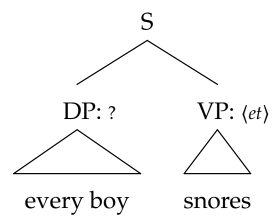

- There are however several reason why an e-type treatment is inadequate. One reason is that such an analysis licenses various invalid inferences.
- For example, consider the following claims.
  i. $\{x \mid x \text{ arrived yesterday morning}\} \subseteq \{x \mid x \text{ arrived yesterday}\}$
  ii. A sentence whose subject denotes an individual x is true iff x is a member of the set denoted by the VP.
- Now suppose that $[\![Bob]\!] \in D_e$. If so, then the following inference is valid.

(122) Bob arrived yesterday morning.

Bob arrived yesterday

- But with certain QNPs, this type of inference is not licensed.

(123) No letter arrived yesterday morning

#No letter arrived yesterday

(124) Exactly one letter arrived yesterday morning

#Exactly one letter arrived yesterday

(125) Less than four letters arrived yesterday morning

#Less than four letters arrived yesterday

- The only viable conclusion is that expressions such as 'no letter', 'exactly one letter', etc. are not e-type expressions.

Another Problem with *e*-Type Denotations

- Consider the following sentence.

(126) Bob is a professor and Bob is not a professor.

- This sentence seems contradictory. We can explain why by making two simple, but plausible, assumptions.

(I.) $\forall A\big(A \subseteq D_e \rightarrow \neg\exists x(x \in A \wedge x \notin A)\big)$

(II.) $[\![Bob]\!] \in D_e$.

- Now compare (126) to (127).

(127) Some woman is a professor and some woman is not a professor.

- (127) is not a contradiction, but if 'some woman' denoted an individual in D_e (perhaps 'the arbitrary woman' as it has been suggested), (127) *should* be a contradiction.
- Hence, we are forced to either abandon principle (I.) above — or give up the assumption that $[\![some\ woman]\!] \in D_e$. It's an easy choice.

8.2.2 QNPs as Type $\langle e,t \rangle$

- Given that QNPs seem to express something about multiple individuals, one might be tempted to treat them as sets of individuals. For example:
  - 'every boy' denotes the following set: $\{x \mid \text{boy}(x)\}$
- Unfortunately, such an anlysis also has some obvious shortcomings. Consider for example what set of individuals the following QNPs should denote?
  - 'no boy' denotes what set? $\{x \mid \neg\text{boy}(x)\}$?
  - 'many boys' denotes what set?
  - 'some boy' denotes what?
- In order to determine the best semantic analysis of QNPs and quantificational determiners, we could try looking at the syntax of sentences containing such expressions.

8.2.3 The Structure of Quantified Sentences in English

- Quantificational sentences in English generally have the following structure.

(128)

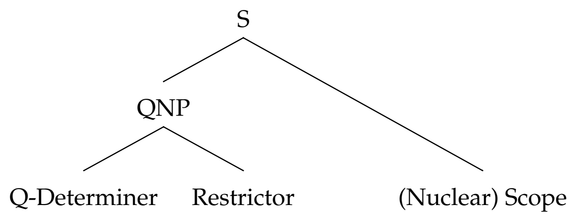

  - A quantificational determiner (e.g. 'every', 'some', 'many', 'most') combines with a predicate (the *restrictor*) to yield a QNP.
  - This QNP then combines with another property (the scope) as its argument and outputs something that has a truth value.

- Here is an example.

(129) Every bird flies.

(130)

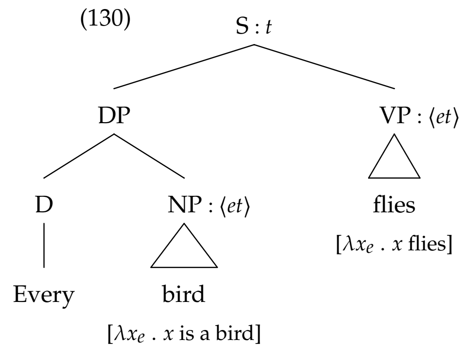

- In order for a type driven derivation of the truth conditions to succeed, the DP-node must combine with the VP-node to yield something of type t.
- The VP-node is of type $\langle et \rangle$ and we have already ruled out that a QNP could denote something of type e or type $\langle et \rangle$.
- Hence, for the derivation to work out here, the DP must be of type $\langle et,t \rangle$.
- And since the quantificational determiner ('every') must combine with something of type $\langle et \rangle$ and output something of type $\langle et,t \rangle$, the semantic type of the quantificational determiner must be $\langle et,\langle et,t \rangle\rangle$.
- So, 'every' should be a function which at least looks something like (131).

(131) $[\![every]\!] = [\lambda P_{\langle et \rangle} \,.\, [\lambda Q_{\langle et \rangle} \,.\, \underline{\hspace{4cm}}]]$

- We know that the value description of the second λ-function must be something of type t, but the crucial question is what it should be.

- Here is a suggestion.

(132) $[\![every]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.$ for every $x \in D_e$ if $P(x) = 1$, then $Q(x) = 1]]$
(133) $[\![every\ bird]\!] = [\lambda Q_{\langle et\rangle} \,.$ for every $x \in D_e$ if $\text{bird}(x)$, then $Q(x) = 1]$

- Equivalently (with a innocuous cheating) in more conspicuous notation.

(134) $[\![every]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.\, P \subseteq Q]]$
(135) $[\![every\ bird]\!] = [\lambda Q_{\langle et\rangle} \,.\, \{x \mid \text{bird}(x)\} \subseteq Q]]$

- This gives us what we want (at least for 'every'). Summing up:
  - A QNP denotes a function from a set of individuals to a truth value — normally this is referred to as a *generalized quantifier* (GQ).
  - A quantificational determiner denotes a function from a set of individuals to a GQ.
  - On this analysis, a quantificational sentence in natural language is now predicted to say something about the relation between two sets of individuals. This type of analysis is thus sometimes called a *relational theory* of determiner denotations.

In generalized quantifier theory, a generalized quantifier is simply a set of sets of individuals, also called a *second-order* property—a property of properties. However, within the confines of our type-theoretic semantics, properties are functions from individuals to truth values, and so a generalized quantifier is a function from sets of individuals to truth values. These are however interdefinable.

- Here are some further examples of quantificational determiner denotations.

(136) $[\![some]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.\, P \cap Q \neq \varnothing\ ]]$

(137) $[\![at\ least\ ten]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.\, |P \cap Q| \geq 10\ ]]$

(138) $[\![exactly\ three]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.\, |P \cap Q| = 3\ ]]$

(139) $[\![no]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.\, |P \cap Q| = 0\ ]]$

- The analysis of natural language QNPs as GQs has a number of significant advantages over the treatment of quantification in FOL where quantification is done using the unary $\forall$ and $\exists$.

Natural Language Quantifiers vs. First Order Quantifiers

- There are several reasons why quantification in natural language cannot be adequately captured by the standard quantifiers of FOL.

1. $\forall$ and $\exists$ are *syncategorematic* expressions. That is, $\forall$ and $\exists$ do not have any meaning in isolation. So, in a formula of the form $\forall x[F(x) \to G(x)]$, no constituent corresponds to 'every' or 'every F'. This violates our compositionality thesis.
2. The meaning of e.g. 'most Fs' cannot be captured in terms of the unary quantifiers of FOL, however defining the meaning of 'most' when using generalized quantifiers is straightforward.

   (140) $[\![most]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.\, |P \cap Q| > |P – Q|\ ]]$
   (141) $[\![most\ Fs]\!] = [\lambda Q_{\langle et\rangle} \,.\, |F \cap Q| > |F – Q|\ ]$

The reason is, in short, that there no way of making a unary quantifier like $\forall$ range over *most* elements of a *subset* of D_e.

8.2.4 Restricted Quantification

- An important property of quantification in natural language which distinguishes it from quantification in FOL is that it is **restricted** quantification. Consider (142a) and its representation in FOL.

(142) a. Some boy snores.
b. $\exists x[\text{boy}(x) \wedge \text{snores}(x)]$

- Intuitively, (142a) says something about only a subset of the individuals in the domain D_e, namely the set of boys. But its first-order representation in (142b) appears to be making a **general** claim about a property of the entire domain of individuals.
- The restricted nature of quantifiers in natural language is particularly evident when one considers what syntactic gerrymandering is needed to capture universal statements in FOL — i.e. notice that (143a) is not intuitively a conditional claim.

(143) a. Every boy snores.
b. $\forall x[\text{boy}(x) \rightarrow \text{snores}(x)]$

- In (143a), a conditional statement is used to mimick restricted quantification. Admittedly, as far as the truth conditions are concerned, (143b) does seem to capture the meaning of (143a).
- Yet, as already mentioned, there is no way of capturing the meaning of many natural language quantificational expressions in terms of the quantifiers of FOL. This means that the strategy of trying to capture all natural language quantification in terms of the quantifiers of FOL is not going to work.
- Another important observation is that quantifiers in natural language are **conservative**, cf. the principle below.

CONSERVATIVITY (CONS)
If Q_{D_e} is a relational quantifier, then Q_{D_e} is conservative iff for each $A, B \subseteq D_e$: $Q_{D_e}(A,B) \Leftrightarrow Q_{D_e}(A, A \cap B)$

- This principle says that for any quantifier Q defined over the domain of D_e of our model $\mathfrak{M}$, if Q is conservative, then the following equivalences hold.

(144) Every student snores $\longleftrightarrow$ Every student is a student who snores.
(145) No student snores $\longleftrightarrow$ No student is a student who snores.
(146) Some student snores $\longleftrightarrow$ Some student is a student who snores.

- This is a significant observation as technically speaking it is immensely easy to formulate (and envision good uses of) various quantificational determiner meanings that are not conservative, e.g. (147) and (148) below.[1]

[1] See Peters and Westerståhl (2007, 138).

(147) $\mathrm{MO}_{D_e}(A,B) \Leftrightarrow |A| > |B|$

(148) $I_{D_e}(A,B) \Leftrightarrow |A| = |B|$
(known as the Härtig quantifier)

- Interestingly, natural languages *in general* just do not appear to have such quantifiers, i.e. quantifiers that violate CONS.

8.3 Directional Entailments and Negative Polarity

- English (but also many other languages) contains a class of expressions called **negative polarity items** (NPIs).
- These are expressions that are only licensed in certain linguistic environments, typically some kind of *negative* environment.
- Examples of NPIs include: 'any', 'ever', 'yet', 'at all', 'budge an inch', 'don't give a hoot'.

(149) a. * John did eat any carrots.
b. John did **not** eat any any carrots.

(150) a. * **Everybody** ate any carrots.
b. **Nobody** ate any carrots.

(151) a. * Lucy **claims** that Peter ever cheated.
b. Lucy **denies** that Peter ever cheated.

(152) a. * **With** anyone signed up, we have no chance of a prize.
b. **Without** anyone signed up, we have no chance of a prize.

(153) a. * Frank arrived **after** anyone else did.
b. Frank arrived **before** anyone else did.

(154) a. * Miriam **usually** budges an inch on issues like this.
b. Miriam **never** budges an inch on issues like this.

- The variations exhibited by just these few examples (there are many more) raises several questions, for example:
  (a) What explains the distribution of NPIs?
  (b) Is the property that licenses NPIs (viz. makes an NPI acceptable in a sentence) syntactic, semantic, or pragmatic?
- One clue towards the answer to this question comes from the theory of GQs. In particular, the *directional entailments* of GQs.
- Consider the following inferences.

(155) Every man snores
$\overline{}$
Every old man snores

(156) Every man snores
$\overline{}$
#Every man snores loudly

(157) Every old man snores
$\overline{}$
#Every man snores

(158) Every man snores loudly
$\overline{}$
Every man snores

- What these examples demonstrate is that the determiner 'every' licenses inferences where:
  a. The NP in the restrictor is replaced with more restrictive NP, for example a replacement of 'man' with 'old man'.
  b. The NP in the scope is replaced with a less restrictive NP, for example a replacement of 'snores loudly' with 'snores'.
- Following Ladusaw (1979), these inferential properties of determiners are referred to as *upwards or downwards directional entailments*.
  - Expressions that license inferences from less to more restrictive terms, cf. (155) are labeled **downwards entailing**.
  - Expressions that license inferences from more to less restrictive terms are labeled **upwards entailing**.
- Formally, these properties are defined as follows:
  - For any quantificational determiner α and any set $A, B \subseteq D$:
    — α is *downwards entailing* in its restrictor term iff $\big(\llbracket\alpha\rrbracket^g(A)(B) \wedge A' \subseteq A\big) \Rightarrow \llbracket\alpha\rrbracket^g(A')(B)$
    — α is *downwards entailing* in its scope term iff $\big(\llbracket\alpha\rrbracket^g(A)(B) \wedge B' \subseteq B\big) \Rightarrow \llbracket\alpha\rrbracket^g(A)(B')$

    — α is *upwards entailing* in its restrictor term iff $\big(\llbracket\alpha\rrbracket^g(A)(B) \wedge A \subseteq A'\big) \Rightarrow \llbracket\alpha\rrbracket^g(A')(B)$
    — α is *upwards entailing* in its scope term iff $\big(\llbracket\alpha\rrbracket^g(A)(B) \wedge B \subseteq B'\big) \Rightarrow \llbracket\alpha\rrbracket^g(A)(B')$
- So, 'every' is downwards entailing in its restrictor, cf. (155), but upwards entailing in its scope, cf. (158).
- However, directional entailment properties vary with various determiners. That is, not every determiner exhibits the same pattern as 'every'. See e.g. the inferences involving 'no' below.

(159) No man snores
$\overline{\text{No old man snores}}$

(160) No man snores
$\overline{\text{No man snores loudly}}$

(161) No old man snores
$\overline{\text{\#No man snores}}$

(162) No man snores loudly
$\overline{\text{\#No man snores}}$

- In contrast to 'every', 'no' is downwards entailing in both its restrictor and scope. Now, compare this to 'some'.

(163) Some man snores
$\overline{\text{\#Some old man snores}}$

(164) Some man snores
$\overline{\text{\#Some man snores loudly}}$

(165) Some old man snores
$\overline{\text{Some man snores}}$

(166) Some man snores loudly
$\overline{\text{Some man snores}}$

- The determiner 'some' is upwards entailing in both its restrictor and scope.

(167) Three men snore
$\overline{\text{\#Three old men snore}}$

(168) Three men snore
$\overline{\text{\#Three men snore loudly}}$

(169) Three old men snore
$\overline{\text{Three men snore}}$

(170) Three men snore loudly
$\overline{\text{Three men snore}}$

- Similar to 'some', 'three' is upwards entailing in both its restrictor and scope.

Directional Entailments and Generalized Quantifiers

- It is important to realize that once quantificational determiners such as 'every', 'no', and 'some' are analyzed in terms of GQs, the directional entailment properties can be trivially predicted on the basis of their semantics.
- For example, let:

$$A = \{x \mid \text{man}(x)\}$$
$$B = \{y \mid \text{snores}(y)\}$$
$$C = \{z \mid \text{old–man}(z)\}$$

- Using generalized quantifiers, we get the following truth conditions:

$$[\![\text{every man snores}]\!]^g = 1 \text{ iff } A \subseteq B$$

— but since $C \subseteq A$, then by transitivity of the $\subseteq$-relation, it follows trivially that if 'every man snores' is true, so is 'every old man snores' too.

Directional Entailment and NPI Licensing

- Now observe the following interesting correlation.

(171) Every politician who has *ever* committed fraud has been convicted.
(172) # Every politician who has committed fraud has *ever* been convicted.

(173) No politician who has *ever* committed fraud has been convicted.
(174) No politician who has committed fraud has *ever* been convicted.

(175) # Some politician who has *ever* committed fraud has been convicted.
(176) # Some politician who has committed fraud has *ever* been convicted.

(177) # Three politicians who have *ever* committed fraud have been convicted.
(178) # Three politicians who have committed fraud have *ever* been convicted.

- Putting this together with our observations about directional entailments yields the following chart.

$\uparrow$ = upwards entailing (set-to-superset) $\downarrow$ = downwards entailing (set-to-subset)

$\checkmark$ = NPIs Licensed

DETERMINER	RESTRICTOR	SCOPE
every	$\downarrow$ $\checkmark$	$\uparrow$
no	$\downarrow$ $\checkmark$	$\downarrow$ $\checkmark$
some	$\uparrow$	$\uparrow$
three	$\uparrow$	$\uparrow$

- Thus, from studying the properties of generalized quantifiers, we are now in a position to draw the following conclusion: negative polarity items (NPIs) are licensed in downwards entailing environments.
- However, is this observation sufficiently general? I.e. can it be extended to cases that do not directly involve determiners? Yes.
- Consider the contrast between the sentences below.

(153a) * Frank arrived **after** <u>anyone</u> else did.
(153b) Frank arrived **before** <u>anyone</u> else did.

- This example demonstrates 'after' does not license NPIs in its scope, but that 'before' does. Yet, this correlates precisely with their intuitive directional entailments. In particular, a preposition such as 'before' is downwards entailing in its scope whereas 'after' is not.

(179) Frank arrived *after* every person had spoken.

#Frank arrived *after* every person had spoken loudly.

(180) Frank arrived *before* every person had spoken.

Frank arrived *before* every person had spoken loudly.

- Consider also temporal adverbs such as 'never' and 'usually'.

(181) * Bertha **usually** lets anyone speak.
(182) Bertha **never** lets anyone speak.

- Again, this difference corresponds directly to a difference in the entailment profiles of the two expressions.

(183) Bertha usually lets Frank speak

#Bertha usually lets Frank speak loudly

(184) Bertha never lets Frank speak

Bertha never lets Frank speak loudly

> **In conclusion**: NPI licensing is the result of a semantic property which is inherent in specific lexical items.

A Problem Case: Conditionals

- NPIs are also licensed in the antecedents of both counterfactual and indicative conditionals, cf. (185) and (186) respectively below.

(185) **If** Jack had passed any exam, his mom would have been proud.
(186) **If** Jack passed any exam, his mom would be proud.

- So, if the above analysis of NPIs is correct, antecedents of conditionals should be downwards entailing.
- The problem is that antecedents of conditionals do not seem to have this property. Specifically, if antecedents of conditionals were downwards entailing, they should license various inferences that are intuitively invalid.
- Since these inferences are licensed when conditionals are analyzed as material implication, this is a well known problem often referred to as the problem of *antecedent strengthening* (or *augmentation)*.
- To illustrate, notice that the following is intuitively an invalid inference.

(187) If Frank had passed the exam, his mom would have been proud.

#If Frank had passed the exam and cheated, his mom would have been proud.

- Yet, if antecedents of conditionals are really downwards entailing, these inferences should sound valid.
- In conclusion, conditionals are a prima facie problem for the analysis of NPIs as expressions that are licensed in downwards entailing environments.

8.4 Interlude: A Sample Derivation Involving a QNP

- **Lexical Entries**
  - $[\![every]\!] = [\lambda P_{\langle et \rangle} \, . \, [\lambda Q_{\langle et \rangle} \, . \, \text{For all } x \in D_e \text{ such that } P(x) = 1, Q(x) = 1]]$
  - $[\![boy]\!] = [\lambda y_e \, . \, y \text{ is a boy}]$
  - $[\![snores]\!] = [\lambda z_e \, . \, z \text{ snores}]$

- By FA

$$\left[\!\!\left[\; [_{\text{DP}} \; [_{\text{D}} \text{ every}] \; [_{\text{NP}} \; [_{\text{N}} \text{ boy}]]] \; \right]\!\!\right] \left(\left[\!\!\left[\; [_{\text{VP}} \; [_{\text{V}} \text{ snores}]] \; \right]\!\!\right] \right)$$

- By NN (twice)

$$\left[\!\!\left[\; [_{\text{DP}} \; [_{\text{D}} \text{ every}] \; [_{\text{NP}} \; [_{\text{N}} \text{ boy}]]] \; \right]\!\!\right] ([\![snores]\!])$$

- By FA

$$\left[\!\!\left[\; [_{\text{D}} \text{ every}] \; \right]\!\!\right] \left(\left[\!\!\left[\; [_{\text{NP}} \; [_{\text{N}} \text{ boy}]] \; \right]\!\!\right] \right) ([\![snores]\!])$$

· By NN (twice)

$$\left[\!\!\left[\begin{array}{c} \text{D} \\ | \\ \text{every} \end{array} \right]\!\!\right] ([\![\textit{boy}]\!])\ ([\![\textit{snores}]\!])$$

· By NN

$$[\![\textit{every}]\!]([\![\textit{boy}]\!])\ ([\![\textit{snores}]\!])$$

· By Lexical Entries

$$\begin{aligned}
& [\![\textit{every}]\!]([\![\textit{boy}]\!])\ ([\![\textit{snores}]\!]) \\
= \ & [\![\textit{every}]\!]([\![\textit{boy}]\!])\ ([\lambda z_e\ .\ z \text{ snores}]) \\
= \ & [\![\textit{every}]\!]([\lambda y_e\ .\ y \text{ boy}])\ ([\lambda z_e\ .\ z \text{ snores}]) \\
= \ & [\lambda P_{\langle et \rangle}\ .\ [\lambda Q_{\langle et \rangle}\ .\ \text{For all } x \in D_e \text{ such that } P(x) = 1, Q(x) = 1]] \\
& ([\lambda y_e\ .\ y \text{ is a boy}])\ ([\lambda z_e\ .\ z \text{ snores}])
\end{aligned}$$

· By β-reduction

$$\begin{aligned}
= \ & [\lambda Q_{\langle et \rangle}\ .\ \text{For all } x \in D_e \text{ such that } [\lambda y_e\ .\ y \text{ is a boy}](x) = 1, Q(x) = 1] \\
& ([\lambda z_e\ .\ z \text{ snores}]) \\
= \ & [\lambda Q_{\langle et \rangle}\ .\ \text{For all } x \in D_e \text{ such that } x \text{ is a boy}, Q(x) = 1]([\lambda z_e\ .\ z \text{ snores}]) \\
= \ & \text{For all } x \in D_e \text{ such that } x \text{ is a boy}, [\lambda z_e\ .\ z \text{ snores}](x) \\
= \ & 1 \text{ iff for all } x \text{ such that } x \text{ is a boy}, x \text{ snores.}
\end{aligned}$$

8.5 Quantifiers in Object Position

- So far, we have only considered sentences where the quantificational expression occurred in subject position. It turns out that a tricky problem arises whenever a QNP occurs in object position.

(188)

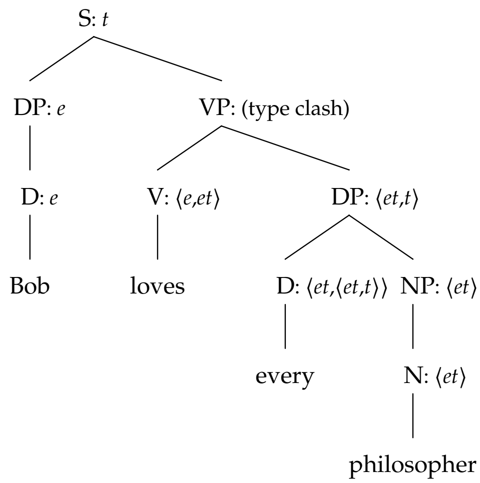

- With two quantifiers, the problem becomes even more dramatic as the sentence now also becomes ambiguous.

(189)

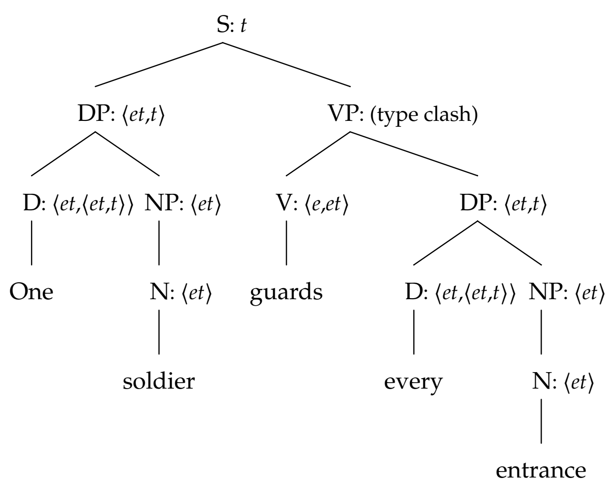

- This sentence has two interpretations that can be roughly paraphrased as follows.

(190) One soldier guards every entrance.

a. There is just one soldier and he guards every entrance
b. For every entrance, one soldier guards that entrance.

8.5.1 *In Situ* Solutions

- One way some people have attempted to solve this problem is by assuming that expressions like 'every F', 'some F' etc. have flexible semantic types.
- On this view, quantificational determiners are systematically ambiguous. Consequently, we need to distinguish between different possible meanings of 'every', in effect two different lexical entries.

(191) $[\![every_1]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle et\rangle} \,.\, \text{for all } x \in D_e \text{ such that } P(x) = 1, Q(x) = 1]]$

(192) $[\![every_2]\!] = [\lambda P_{\langle et\rangle} \,.\, [\lambda Q_{\langle e,et\rangle} \,.\, [\lambda x_e \,.\, \text{for all } y \in D_e$
$\text{such that } P(y) = 1, Q(y)(x) = 1]]]$

- (192) is the semantic type required for when a quantifier occurs in direct object position—the semantic type of this expression is $\langle et, \langle\langle e,et\rangle, et\rangle\rangle$
- This can be generalized to all determiners and expanded to include denotations for cases where quantifiers are in object or indirect object position of a ditransitive verb like 'gives'.
- This solution might seem somewhat ad hoc, but there several reasons for thinking that certain lexical items must have flexible types (and so multiple lexical entries). Consider the following sentences.

(193) $[_{\text{S}}$ $[_{\text{S}}$ Bob bought a car] and $[_{\text{S}}$ Mary bought a car]]
(194) $[_{\text{S}}$ $[_{\text{DP}}$ $[_{\text{DP}}$ Bob] and $[_{\text{DP}}$ Mary]] bought a car]
(195) $[_{\text{S}}$ Mary $[_{\text{VP}}$ $[_{\text{VP}}$ bought a car] and $[_{\text{VP}}$ sold a boat]]]
(196) $[_{\text{S}}$ Bob works $[_{\text{PP}}$ $[_{\text{PP}}$ in the school] and $[_{\text{PP}}$ in the front office]]]

- In (193), where the conjunction is coordinating two sentences, the standard analysis of 'and' as a **sentential** connective (as a function from truth values to a function from truth values to truth values $\langle t,\langle t,t\rangle\rangle$) makes the right predictions.
- However, in (194) the conjunction is coordinating DPs, in (195) the conjunction is coordinating VPs, and in (196) the conjunction is coordinating PPs. So, we need flexible types for 'and' (and the same for 'or').

CONSTITUENT	COORDINATION	SEMANTIC TYPE
S	S *and* S	$\langle t,\langle t,t\rangle\rangle$
VP	VP *and* VP	$\langle et,\langle et,et\rangle\rangle$
PP	PP *and* PP	$\langle et,\langle et,et\rangle\rangle$
DP_e	DP_e *and* DP_e	$\langle e,\langle e,e\rangle\rangle$

- But ambiguity solutions are never happy solutions.

8.5.2 Movement Solutions

- An alternative solution is to assume that the phrase structures for sentences containing quantifiers in object position actually look significantly different from their surface representations.
- As mentioned earlier, it is a central assumption in much contemporary syntactic theory that there is an (extra) layer of representation called *Logical Form* (LF). The structure of an LF can diverge significantly from the structure of a representation of the surface structure.
- In LFs, quantifiers will have been moved by application of a syntactic rule called **Quantifier Raising** (QR).

Actually, in syntax the QR rule has now been replaced with a more general rule of movement called "move α" which can be applied to other constituents than quantifier phrases.

- For example, the LF of (197) would look something like (198). In this representation, an application of QR has resulted in the movement of the quantifier phrase 'every philosopher'. This phrase is then **adjoined to** *S* leaving behind a trace in its originating position.

(197) Bob loves every philosopher.

(198)

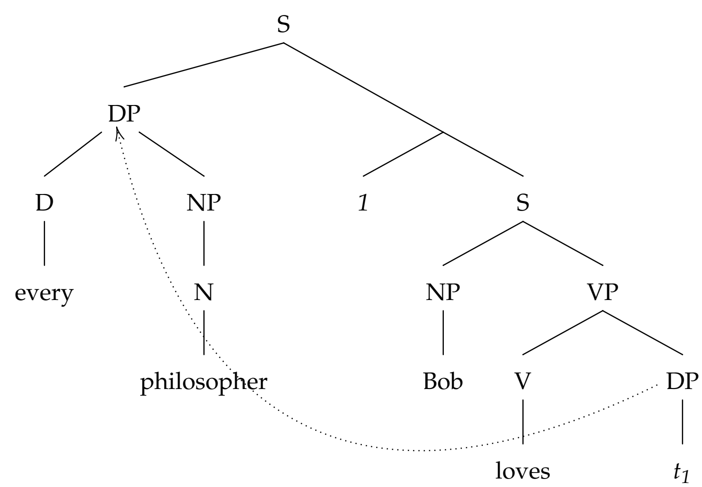

8.5.3 Interlude: Evidence for Movement?

- From the point of view of semantics, the assumption that there are LFs and that these are the result of the application of various transformational rules (e.g. QR) might seem odd (and ad hoc).

  However, many syntactic arguments have been put forward in favor of this claim. Here is a brief outline of a couple of these arguments.

Wh-movement

- English is a SVO language (subject, verb, object) as witnessed by simple declarative sentence such as (199) and the question in (200).

(199) Bob admires Mary.
(200) Bob admires which philosopher?

- Notice that *wh*-phrases normally occur at the beginning of a sentence. Because of this, it has been hypothesized that *wh*-phrases are displaced — they actually undergo movement from their original position and adjoin to S, cf. e.g. (201).

(201) Which philosopher does Bob admire?
(202) [Which philosopher]$_1$ (does) Bob admire t_1?

Actually, the argument here is not so strong in English as this kind of case marking does not seem mandatory anymore in English. However, the same observation can be made for e.g. German where correct case marking remains absolutely mandatory.

- This hypothesis is supported by e.g. observations concerning case marking in embedded questions.

(203) a. Bob knows whom Mary met.
b. * Bob knows who Mary met. (* = ungrammatical)

- The most natural explanation for the ungrammaticality of (203b) is that the relative pronoun ('whom') originally served as the direct object of the sentence, but that it has moved and been adjoined to the embedded S.

Quantifier Movement

- You might be wondering why evidence for *wh*-phrase movement should give us any reason to think that QNPs also move?
- **A Preliminary Consideration**
  - Quantifiers and certain *wh*-phrases are both syntactically the result of combining a determiner with a noun phrase, e.g. 'every philosopher' — 'which philosopher'. In other words, quantifiers and *wh*-phrases are both DPs.
- **Quantifier Scope**
  We have already seen examples of sentences containing two quantifier phrases and that this leads to ambiguity.
  - This is a systematic phenomenon. Sentences containing two QNP typically have at least two possible interpretations.
  - These ambiguities are often assumed to be the result of differing quantifier scopes. For example, in FOL, one distinguishes between the following.

    (204) One soldier guards every entrance.
    a. $\exists x[\text{soldier}(x) \wedge \forall y[\text{entrance}(y) \rightarrow x \text{ guards } y]]$ $(\exists > \forall)$
    b. $\forall x[\text{entrance}(x) \rightarrow \exists y[\text{soldier}(y) \wedge y \text{ guards } x]]$ $(\forall > \exists)$

  - It seems natural to assume that quantifier ambiguities in natural language should also be captured in terms of relative scope. If QNP can move by applications of QR that might explain why a QNP that takes narrow scope on the surface can nevertheless be interpreted as taking wide scope.

- **Further Evidence: Movement and Coordination**
  Another important similarity between *wh*-phrases and quantifiers is that their movement appears to be prohibited in the same syntactic environments.

  - **Coordinate Structure Islands**
    Consider the sentences below.

    (205) Bob saw Chomsky and which philosopher?
    (206) * Which philosopher did Bob see Chomsky and?

  - We see the same pattern in (207) below.

    (207) Some linguist saw Chomsky and every philosopher.
    a. $\exists x[\text{linguist}(x) \wedge \text{saw}(x, \text{Chomsky}) \wedge \forall y[\text{philosopher}(y) \rightarrow \text{saw}(x,y)]]$
    b. * $\forall x[\text{philosopher}(x) \rightarrow \exists y[\text{linguist}(y) \wedge \text{saw}(y, \text{Chomsky}) \wedge \text{saw}(y,x)]]$

  - If the quantifier 'every philosopher' could move, then the interpretation in (207b) should be available, i.e. an interpretation where for every philosopher x, there is a linguist who saw Chomsky and x.
  - But that interpretation is not available for (207). This is evidence that the quantifier is prohibited from moving out of a coordinated conjunction (just like a *wh*-phrase).

- There are many other arguments for QR than those discussed above.

Traces

- What about the assumption that QNPs or *wh*-phrases that have been moved and adjoined to S leave behind a trace? What is the justification for this assumption?

  - Consider the following questions.

    (208) Which philosopher does Bob want to invite to the party?
    (209) Which philosopher does Bob want to invite Sue to the party?

  - Now, consider what happens if we contract on 'want' and 'to'.

    (210) Which philosopher does Bob wanna invite to the party?
    (211) * Which philosopher does Bob wanna invite Sue to the party?

  - In (210), the contraction is acceptable, but in (211) it leads to ungrammaticality.
  - This can be explained if it is assumed that the *wh*-phrase 'which philosopher' has been displaced from its original position.

    (212) a. Bob wants to invite **which philosopher** to the party?
    b. [which philosopher]$_1$ (does) Bob want to invite t_1 to the party?

    (213) a. Bob wants **which philosopher** to invite Sue to the party?

b. [which philosopher]$_1$ (does) Bob want t_1 to invite Sue to the party?

- Notice that in (213b), the trace sits in between 'want' and 'to'. If so, this straightforwardly explains why the contraction to 'wanna' is unacceptable. There is a syntactic constituent blocking it.
- The same point is illustrated by the *wh*-question in (214) which is ambiguous, cf. the possible disambiguations below.

(214) Who do you want to succeed?
a. You want who to succeed?
b. You want to succeed who?

- Notice that (215) has only one possible interpretation, namely the one paraphrased in (214a).

(215) Who do you wanna succeed?

- Again, the explanation is that a trace is blocking a contraction of 'want' and 'to'.

(216) a. Who$_1$ (do) you want t_1 to succeed?
b. Who$_1$ (do) you want to succeed t_1?

8.6 Back to Quantifier Phrases in Object Position

- Given the arguments above, we are going to pursue a solution to the problem of quantifiers in object position that relies on quantifier movement, viz. applications of QR.
- We thus assume that the LF for (197) is the phrase structure in (198).

(197) Bob loves every philosopher.

(198)

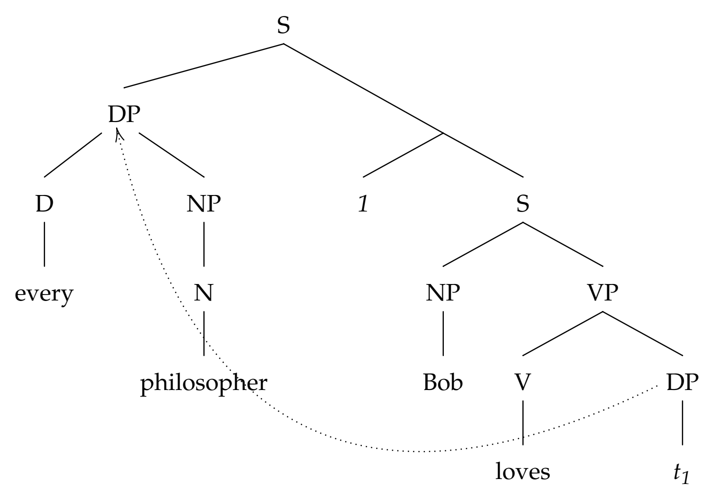

- There is one crucial assumption here, namely that when a quantifier moves across S and adjoins to a higher S, a binder for the trace t_i is created in that process. This binder i then serves the same purpose as a relative pronoun in a restrictive relative clause.
- The result is a structure that can be interpreted using (a slightly modified version of) predicate abstraction.

> **Predicate Abstraction Rule (PA)** (*generalized*)
> Let α be a branching node with daughters β and γ, where β dominates only a numerical index i. Then for any variable assignment g: $[\![\alpha]\!]^g = \lambda x_e \,.\, [\![\gamma]\!]^{g^{[x/i]}}$

Generalized Predicate Abstraction Rule

- I.e. the following is going to happen when we interpret the right branch of (198).

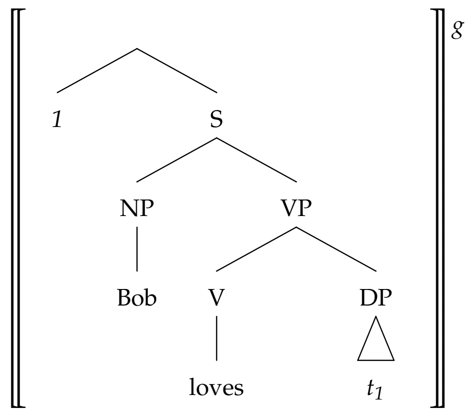

· By Predicate Abstraction

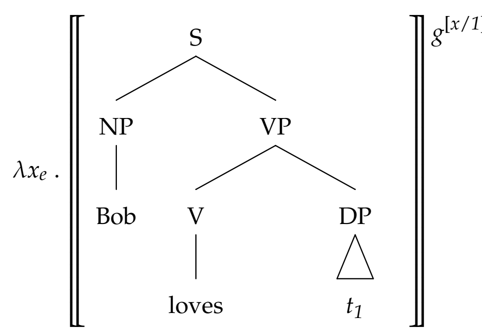

- The restrictive relative clause is now going to end up denoting the property $[\lambda x_e \,.\,$ Bob loves $x]$. And so...

$$[\![every]\!]([\![philosopher]\!])([\lambda x_e \,.\, \text{Bob loves } x])$$
$$= [\![every]\!]([\lambda y_e \,.\, y \text{ is a philosopher}])([\lambda x_e \,.\, \text{Bob loves } x])$$

$= [\lambda P_e \,.\, [\lambda Q_e \,.$ for all $z \in D_e$ such that $P(z) = 1, Q(z) = 1]]$
$([\lambda y_e \,.\, y$ is a philosopher$])([\lambda x_e \,.$ Bob loves $x])$

$= [\lambda Q_e \,.$ for all $z \in D_e$ such that $[\lambda y_e \,.\, y$ is a philosopher$](z) = 1, Q(z) = 1]$
$([\lambda x_e \,.$ Bob loves $x])$

$= [\lambda Q_e \,.$ for all $z \in D_e$ such that z is a philosopher, $Q(z) = 1]([\lambda x_e \,.$ Bob loves $x])$

$=$ For all $z \in D_e$ such that z is a philosopher, $[\lambda x_e \,.$ Bob loves $x](z) = 1$

$=$ 1 iff for all philosophers z, Bob loves z

8.7 Arguments for Movement Solutions

8.7.1 Scope Ambiguities and Inverse Scope

- As pointed out earlier, sentences with two (or more) quantificational expressions are ambiguous.
- However, if QNPs are interpreted *in situ* (i.e. using flexible types), it is unclear how one is supposed to predict that a sentence such as (217) has more than one interpretation.
- With the a rule like QR, different LFs for can be generated for sentences containing multiple QNPs and scope ambiguities can thus be straightforwardly captured. For example, for a sentence such as (217), there are (at least) two associated LFs.

(217) Every boy likes a girl.

a.

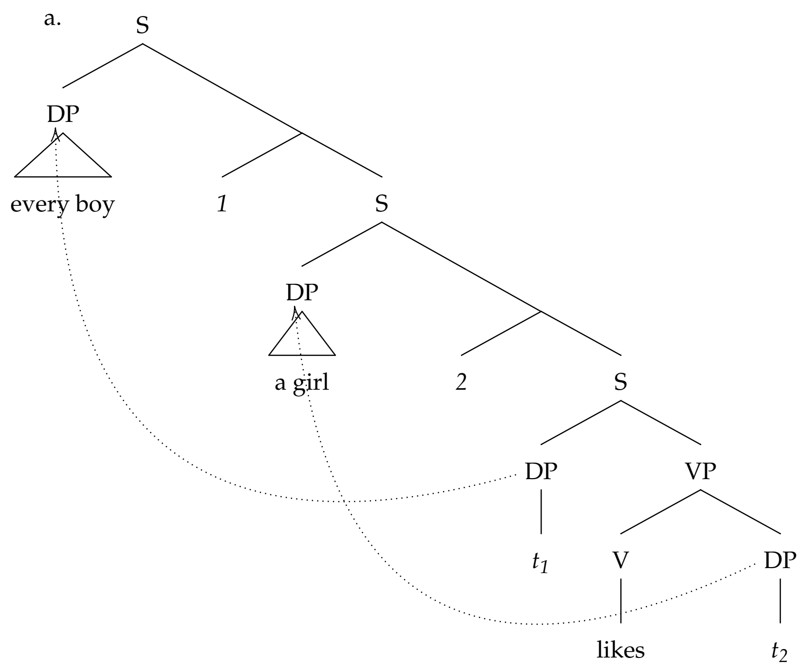

(217) b.

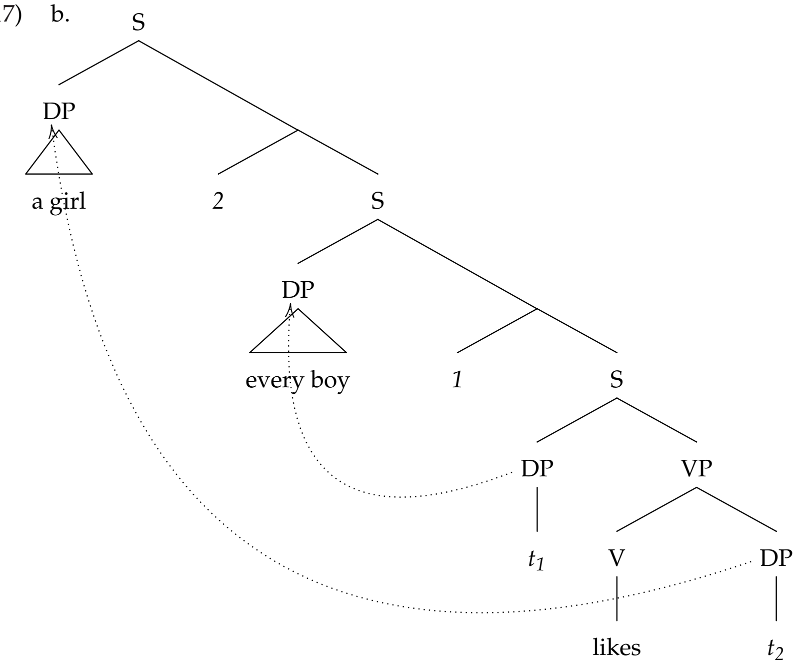

- Both of these LFs are interpretable using our current rules. (217a) yields the wide scope existential interpretation and (217b) yields the narrow scope existential interpretation—as needed.
- Notice also that QNPs are sometimes embedded inside other DPs.

(218) An apple in every basket is rotten
[$_S$ [$_{DP}$ an apple [$_{PP}$ in [$_{DP}$ every basket]]] [$_{VP}$ is rotten]] (surface structure)

- (218) only has one natural reading and on the movement approach this reading is easily captured by allowing the quantifier to move out of the PP and adjoin to S, viz. (219).

(219) [$_S$ [$_{DP}$ every apple] *1* [$_S$ [$_{DP}$ one apple in t_1] [$_{VP}$ is rotten]]] (LF)

8.7.2 Antecedent Contained Deletion

- Consider the sentences below.

(220) Russell wrote *Principia* and Moore did too.
Russell wrote *Principia* and Moore [wrote *Principia*] too.

(221) Schönfinkel discovered a proof before Curry.
Schönfinkel discovered a proof before Curry [discovered a proof].

- In (220) and (221), a verb phrase has been deleted on the surface—this is generally known as VP-ellipsis (verb phrase ellipsis).
    - The general hypothesis about VP-ellipsis is that in the syntactic derivation from SS (surface structure) to PF (phonetic form), the VP is deleted, i.e. it is unpronounced — yet the VP is not deleted in the derivation of the LF.
    - So, the (phonetically) deleted VP is syntactically represented at the level of LF.
    - It is generally assumed that VP-deletion is licensed only when there is a preceding VP that can be copied. In other words, a VP cannot simply be omitted if there is no antecedent from which it could be copied.

There are several other types of deletion/ellipsis in natural language, for example (a) *noun phrase deletion* as in "I enjoy reading Irene's papers, but I don't enjoy reading Richard's", (b) *sluices* as in "Bob published a book, but I don't know which", and (c) *gaps* as in "Bob married Liz, and Jack Kate".

- Now consider the sentence below.

(222) Church proved every theorem that Fitch did.

- This an example of what is called *antecedent contained deletion* (ACD)—the deleted material is *contained* in the antecedent. The problem with ACDs is that if the deleted material (the VP) is simply copied over, this triggers an infinite regress.

(223) Church proved every theorem that Fitch [proved every theorem that Fitch ... [proved every theorem that Fitch ...] ...] ...]

- Hence, what has to be deleted in (222) is not the surface VP 'proved every theorem that ...', but rather a VP containing a verb and a trace, viz. 'proved t_1'. But the surface structure of (222) does not provide such an antecedent.
- What is needed is an LF where the QNP 'every theorem' has been adjoined to S and left behind a trace, i.e. an LF roughly like (224).

(224)

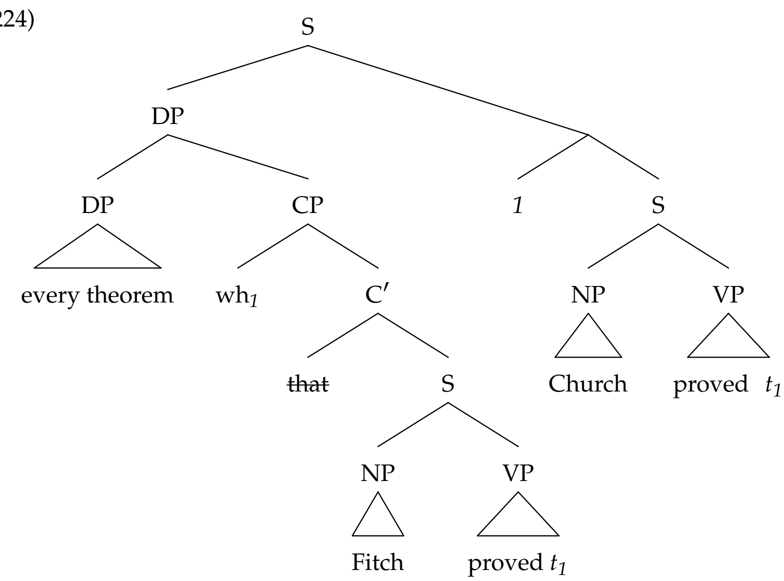

- In other words, the assumption that quantifiers move by QR helps provide an elegant explanation of ACD.

8.7.3 Bound Variable Anaphora

- Quantifiers can *bind* pronouns, as in (225)–(227) below.

(225) Jason loves himself.
(226) No philosopher loves himself.
(227) Every philosopher loves himself.

- The pronouns in (225)–(227) are *reflexive* pronouns and these are necessarily *anaphoric*—viz. acceptable only when a *linguistic* antecedent is available.
- The main characteristic of an *anaphoric* pronoun is that its semantic value is determined by its linguistic antecedent. In (225), since the linguistic antecedent for 'himself' is 'Jason' and the name 'Jason' *refers* to Jason, the reflexive pronoun also refers to Jason — so the sentence means that Jason loves Jason.
- Non-reflexive pronouns have both anaphoric and non-anaphoric uses, cf. (228a) and (228b) respectively.

(228) No philosopher wanted Chomsky to ask him a question.
   a. [No philosopher]$_1$ wanted Chomsky to ask him$_1$ a question.
   b. [No philosopher]$_1$ wanted Chomsky to ask him$_2$ a question.

- The question is what the semantic value of 'himself' and 'him' is when the linguistic antecedent is a QNP.
- If we allow that quantifiers can move even when there is no risk of a type mismatch, we can straightforwardly treat these pronouns as bound variables.
- That is, from (226), we can derive the LF in (229).

(226) No philosopher loves himself.
(229)

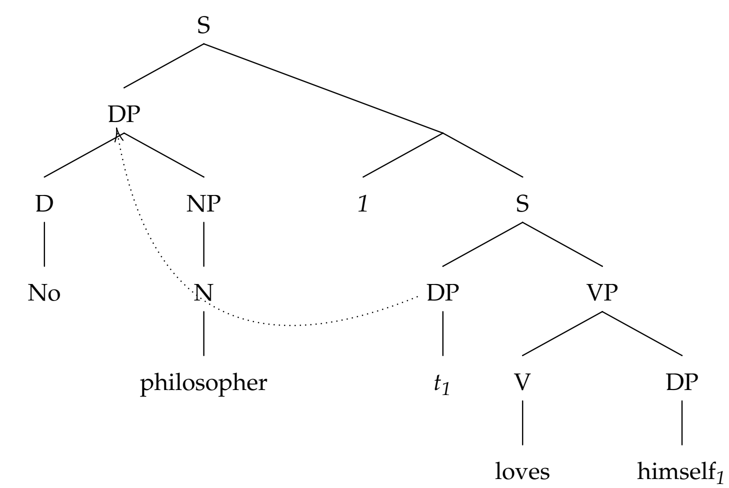

- On the *in situ* approach, getting the truth conditions right for (226) looks problematic.
- In particular, if we assume that quantifiers do not move, the relevant phrase structure for (226) is (230) below.

(230)

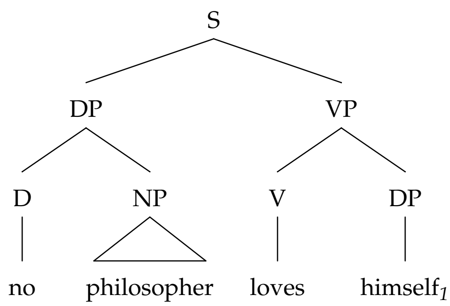

- The obvious problem here is there is no binder for the pronoun 'himself'.
- This means that the meaning of (230) ends up depending on what the variable assignment g maps the index *1* to—i.e.

If	$g(1)$ = Bob	*then*	$[\![(230)]\!]^g = 1$ iff no philosopher loves Bob.
If	$g(1)$ = David	*then*	$[\![(230)]\!]^g = 1$ iff no philosopher loves David.
If	$g(1)$ = Alonzo	*then*	$[\![(230)]\!]^g = 1$ iff no philosopher loves Alonzo.
If	$g(1)$ = Bertrand	*then*	$[\![(230)]\!]^g = 1$ iff no philosopher loves Bertrand.

- These predictions are obviously incorrect, but there is no easy fix for this.
- Even if proponents of an *in situ* analysis assumed that the LF for (230) contained a binder, as below, this fails to solve the problem.

Notice that in order for the VP to denote the right property, namely the following: $[\lambda x_e \,.\, x$ loves $x]$, two identical arguments are needed. A binder for the pronoun only provides one such argument—in other words, the trace makes all the difference here.

(231)

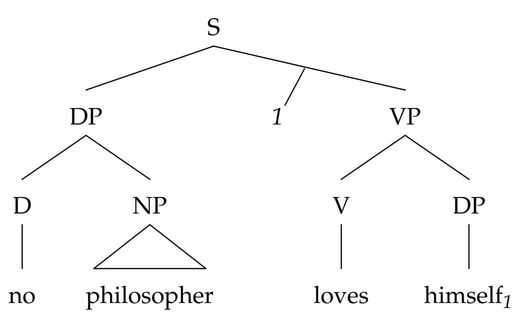

- To deal with this problem, one would need to introduce new (and fairly ad hoc) rules. So, the movement solution seems to have the advantage.
- However, there are various sophisticated analyses out there that deal with binding without assuming anything like LFs or QR, see e.g. Jacobson (2014).

8.8 Derivation: "Every boy loves a girl"

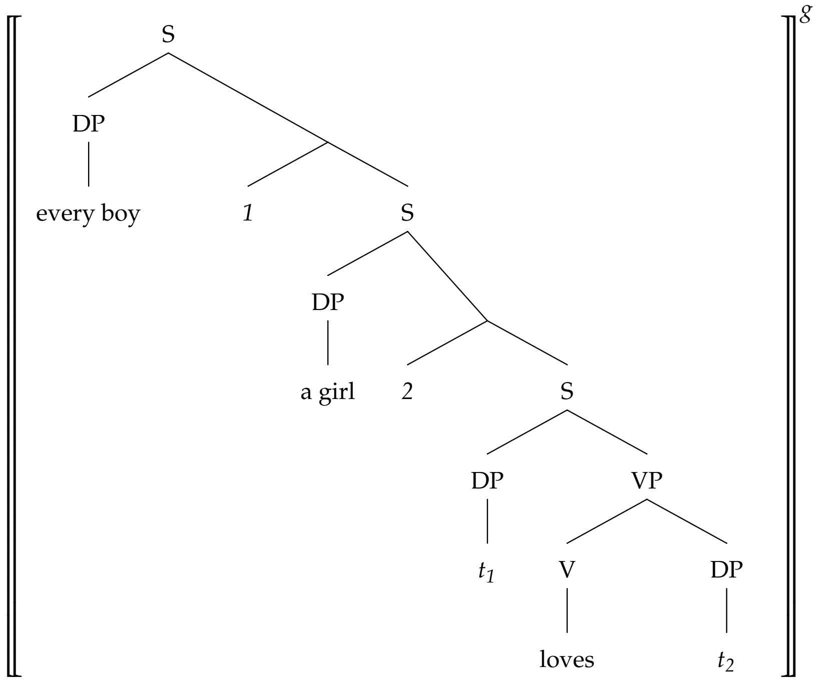

Step 1: FA

$[\![\textit{every boy}]\!]^g$

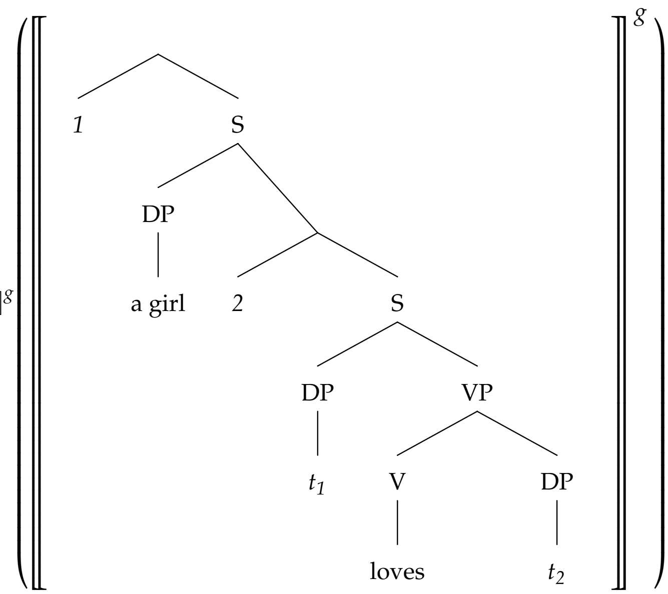

Step 2: PA

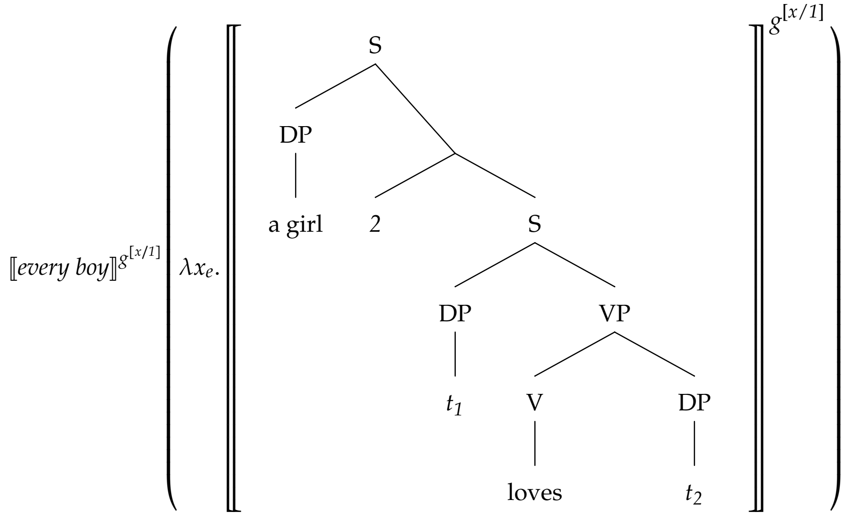

Step 3: FA

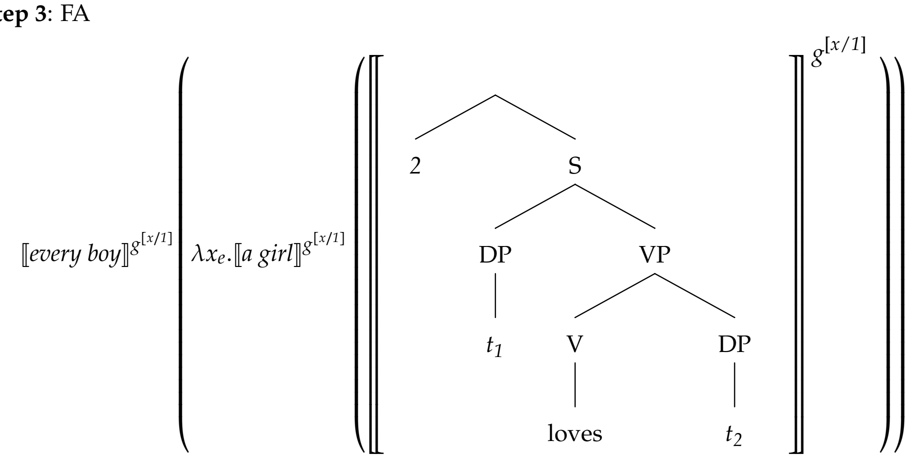

Step 4: FA

$$[\![\textit{every boy}]\!]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\left(\lambda x_e.[\![\textit{a girl}]\!]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\left(\lambda y_e.\left[\!\!\left[\;[_{\text{S}}\;[_{\text{DP}}\;t_1]\;[_{\text{VP}}\;[_{\text{V}}\;\text{loves}]\;[_{\text{DP}}\;t_2]]]\;\right]\!\!\right]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\right)\right)$$

Step 5,6: FA + 1 application of NN to DP-node

$$[\![\textit{every boy}]\!]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\left(\lambda x_e.[\![\textit{a girl}]\!]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\left(\lambda y_e.\left[\!\!\left[\;[_{\text{VP}}\;[_{\text{V}}\;\text{loves}]\;[_{\text{DP}}\;t_2]]\;\right]\!\!\right]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\left([\![t_1]\!]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\right)\right)\right)$$

Step 7: Pronouns and Traces Rule

$$[\![\textit{every boy}]\!]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\left(\lambda x_e.[\![\textit{a girl}]\!]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}\left(\lambda y_e.\left[\!\!\left[\;[_{\text{VP}}\;[_{\text{V}}\;\text{loves}]\;[_{\text{DP}}\;t_2]]\;\right]\!\!\right]^{g\left[\begin{smallmatrix} x/1 \\ y/2 \end{smallmatrix}\right]}(x)\right)\right)$$

Step 8,9,10: FA + 2 applications of NN to V-node and DP-node

$$[\![every\ boy]\!]^{g\begin{bmatrix}x/1\\y/2\end{bmatrix}}\left(\lambda x_e.[\![a\ girl]\!]^{g\begin{bmatrix}x/1\\y/2\end{bmatrix}}\left(\lambda y_e.[\![loves]\!]^{g\begin{bmatrix}x/1\\y/2\end{bmatrix}}([\![t_2]\!]^{g\begin{bmatrix}x/1\\y/2\end{bmatrix}})(x)\right)\right)$$

Step 11: Pronouns and Traces Rule

$$[\![every\ boy]\!]^{g\begin{bmatrix}x/1\\y/2\end{bmatrix}}\left(\lambda x_e.[\![a\ girl]\!]^{g\begin{bmatrix}x/1\\y/2\end{bmatrix}}\left(\lambda y_e.[\![loves]\!]^{g\begin{bmatrix}x/1\\y/2\end{bmatrix}}(y)(x)\right)\right)$$

Step 12,13,14: 3 Applications of Assignment Independent Denotations

$[\![every\ boy]\!](\lambda x_e\ .\ [\![a\ girl]\!](\lambda y_e\ .\ [\![loves]\!](y)(x)))$

Step 15: TN (from the lexicon)

$[\![every\ boy]\!](\lambda x_e\ .\ [\![a\ girl]\!](\lambda y_e\ .\ [\lambda z_e\ .\ [\lambda r_e\ .\ r \text{ loves } z]](y)(x)))$

Step 16: β-reduction

$[\![every\ boy]\!](\lambda x_e\ .\ [\![a\ girl]\!](\lambda y_e\ .\ [\lambda r_e\ .\ r \text{ loves } y](x)))$

Step 17: β-reduction

$[\![every\ boy]\!](\lambda x_e\ .\ [\![a\ girl]\!]([\lambda y_e\ .\ x \text{ loves } y]))$

Step 18: TN (from the lexicon)

$[\![every\ boy]\!](\lambda x_e\ .\ [\lambda P_{\langle et\rangle}\ .\ \text{For some girl } z, P(z) = 1]([\lambda y_e\ .\ x \text{ loves } y]))$

Step 19: β-reduction

$[\![every\ boy]\!](\lambda x_e\ .\ \text{For some girl } z, [\lambda y_e\ .\ x \text{ loves } y](z) = 1)$

Step 20: β-reduction

$[\![every\ boy]\!]([\lambda x_e\ .\ \text{For some girl } z, x \text{ loves } z])$

Step 21: TN (from the lexicon)

$[\lambda P_{\langle et\rangle}\ .\ \text{For every boy } y, P(y) = 1]([\lambda x_e\ .\ \text{For some girl } z, x \text{ loves } z])$

Step 22: TN (from the lexicon)

For every boy y, $[\lambda x_e\ .\ \text{For some girl } z, x \text{ loves } z](y)$

= 1 iff for every boy y, there is some girl z such that y loves z

QED

8.9 Final Notes on Quantifier Movement

- In our solution to the problem of QNPs in object position, we appealed to the rule of quantifier movement.

8.9.1 Quantifier Adjunction

- We generally assumed that raised QNPs are adjoined to S. However, there are strong reasons to believe that raised quantifiers must sometimes adjoin to e.g. VPs, DPs, and PPs.
- The reasons are complex and beyond the scope of this course, but see Heim and Kratzer (1998, 215-235) for discussion.

8.9.2 Movement (Island) Constraints

- As mentioned earlier, quantifier and *wh*-movement is subject to various restrictions. We looked at only one example, namely raising QNPs out of coordinated conjunctions. This constraint is actually more general as it applies to all coordinated environments (and hence was aptly called *the coordinate structure constraint* by Ross (1967).
- Yet, coordinated structures are just one among many syntactic environments which prohibit movement (sometimes called *scope islands*). One environment that deserves mentioning here is relative clauses.
- Consider the sentence below. This sentence contains two quantifier phrases so we should expect it to be ambiguous.

(232) One general who every soldier feared forgot to sound the alarm.

- I hope it is obvious that (269) does not in fact have a wide-scope universal interpretation, i.e. the interpretation paraphrased below.

(233) For every soldier, there was one feared general who forgot to sound the alarm.

- We explain the fact that this reading is unavailable by observing that restrictive relative clauses are scope islands—i.e. one cannot raise a QNP out of a relative clause to take wide scope.

Actually, it's not quite this simple. Here the question of whether QNPs must adjoin to S is relevant, because it seems that for this structure to be interpretable, the quantifier must be raised—only CP-internally. And, importantly, the island constraint described here only prohibits raising the quantifier out of CP to take scope over the QNP 'one soldier'.

Chapter 9

Interlude: Kaplan on Indexicals

9.1 Indexicals and Demonstratives

- One of the most influential papers in philosophy of language and formal semantics is David Kaplan's "Demonstratives" which introduced a semantic framework and analysis of what Kaplan called *indexicals* and *demonstratives*.
- This class of words include pronominal expressions such as 'I', 'you', 'he', 'his', 'she', 'it', 'that', 'this' and adverbs such as 'here', 'now', 'tomorrow', 'today', 'actual', 'present', etc.
- Some of the general questions that Kaplan addresses in "Demonstratives" include:

  I. How to capture within a formal semantics that indexicals are words that have *unstable* meanings — i.e. that indexicals are context sensitive.

  II. How to capture that indexicals are not mere placeholders for objects, i.e. that each individual indexical also seems to have a *stable* meaning that distinguishes it from other indexicals.

  III. How to explicate the role of the context in determining semantic values, i.e. what it means to say that context determines the meaning of indexicals?

- **Unifying Feature of Indexicals**: Each individual indexical "provides a rule which determines the referent in terms of certain aspects of the context" (Kaplan, 1989, 490)
- **A Distinction**
  - **True Demonstratives**: 'that', 'this', 'those', 'they', 'he', 'she', 'you' etc.
    Demonstratives are indexicals that require an associated demonstration to be complete. It is part of the linguistic rules of use that when a speaker uses these words, the speaker must demonstrate the individual/object to which she intends to refer.
  - **Pure Indexicals**: 'I', 'my', 'here', 'now', 'today', etc.
    Pure indexicals are indexicals that rely on context for their meaning, but require no accompanying demonstration.

9.2 Two Seemingly Inconsistent Principles

i. The referent of a pure indexical depends on the context and the referent of a (true) demonstrative depends on the associated demonstration. So, the meaning (the semantic value) of both types of expressions vary with context.

   So, if α is an indexical, it is possible that $[\![\alpha]\!]^c \neq [\![\alpha]\!]^{c'}$

ii. Indexicals are *directly referential* and *rigid*.

- What does it mean for an expression to be *directly referential* and *rigid*?

  - **Rigidity**
    An expression α is rigid iff the referent of α, once determined, is fixed for all possible circumstances (i.e. all possible worlds).

    So, an expression α is rigid iff once α is fixed at w, then $\forall w'\big([\![\alpha]\!]^w = [\![\alpha]\!]^{w'}\big)$.

  - **Direct Reference**
    An expression α is directly referential iff its referent is determined directly and so without the mediation of any descriptive propositional constituent.

- However, note:

  > [...] in interpreting the phrase 'designates the same object in all circumstances' [...] we do not mean that the expression *could not have been used* to designate a different object. We mean rather that given a use of the expression, we may ask of *what has been said* whether *it* would have been true or false in various counterfactual circumstance, and in such counterfactual circumstances, which are the individuals relevant to determining truth-value. (Kaplan, 1989, 493-494)

- **Example I.**
  If Caroline asserts (234) and I assert (235) while pointing to Caroline, intuitively we have said the same thing.

(234) I am a linguist at the University of Edinburgh.
(235) You are a linguist at the University of Edinburgh.

- **Example II.**
  If I assert (236) pointing to Caroline, what I have said is (237).

(236) She is a linguist at the University of Edinburgh.
(237) Caroline is a linguist at the University of Edinburgh.

9.3 Context vs. Circumstance of Evaluation

- This brings us to a crucial distinction in Kaplan's paper — the distinction between CONTEXT of utterance and CIRCUMSTANCE of evaluation.

- On Kaplan's view, determining the meaning of a sentence is a two-stage process. First the sentence is paired with a context to determine 'what is said' or the proposition expressed. Then, to assess whether the sentence is true/false, it has to be paired with a world parameter.

Later, we will add further parameters such as a *time* parameter and a *location* parameter

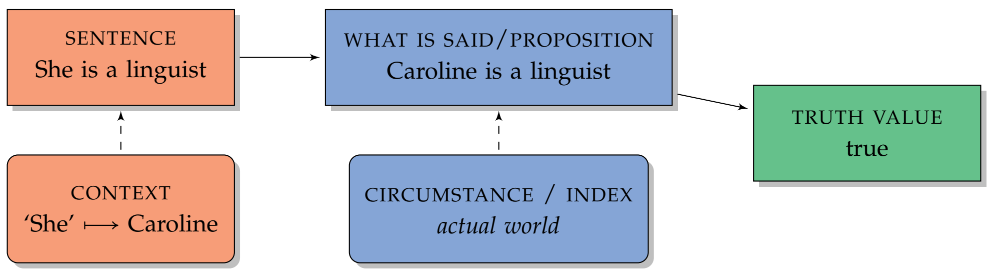

- Given this picture of sentence interpretation, it should now be clearer what Kaplan means when he writes that:

  a. **The referent of an indexical depends on the context.**

     "A directly referential term *may* designate different objects when used in different *contexts*" (1989, 494)

  b. **Indexicals are rigid and directly referential.**

     "But when evaluating what was said in a given context, only a single object will be relevant to the evaluation in all circumstances" (1989, 494)

The Kaplanian Picture (continued)

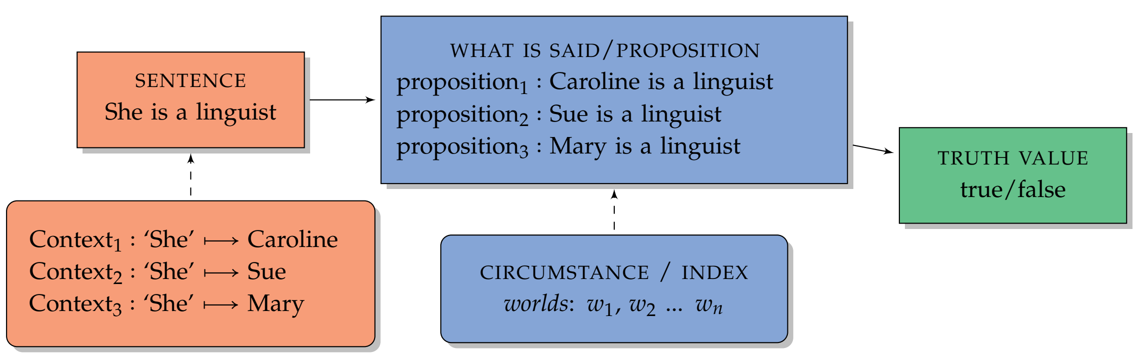

- So, in other words, on Kaplan's view, sentence meaning is relativized along two axes, namely CONTEXT and CIRCUMSTANCE.
- And this should be reflected in our interpretation function: $[\![\cdot]\!]^{c,w}$

9.3.1 Direct Reference vs. Rigid Designators

- Notice that when an expression is directly referential, the relation between the expression and its referent is direct, i.e. the semantic value of the term is not mediated or determined by a propositional constituent (i.e. by a constituent of the proposition expressed).
- What is contributed to the content/proposition expressed is simply the object. As Kaplan puts it:

> [...] the descriptive meaning of a directly referential term is not part of the propositional content (Kaplan, 1989, 497)

- However, an expression can be rigid without being directly referential, for example:

(238) The number n such that:
(Roses are red and $n^2 = 9$) or (Roses are not red and $2^2 = n + 1$)

- This description will denote the same number, 3, in all possible worlds, but this is something that in a sense has to be *determined* at each possible world. The semantic value of the description is a function of a propositional constituent — it just so happens that it denotes the same object in every world.
- Hence the expression above is rigid, but it is not directly referential.

9.3.2 Determining Reference

- The difference between *indexicals*, *names*, and *descriptions*:

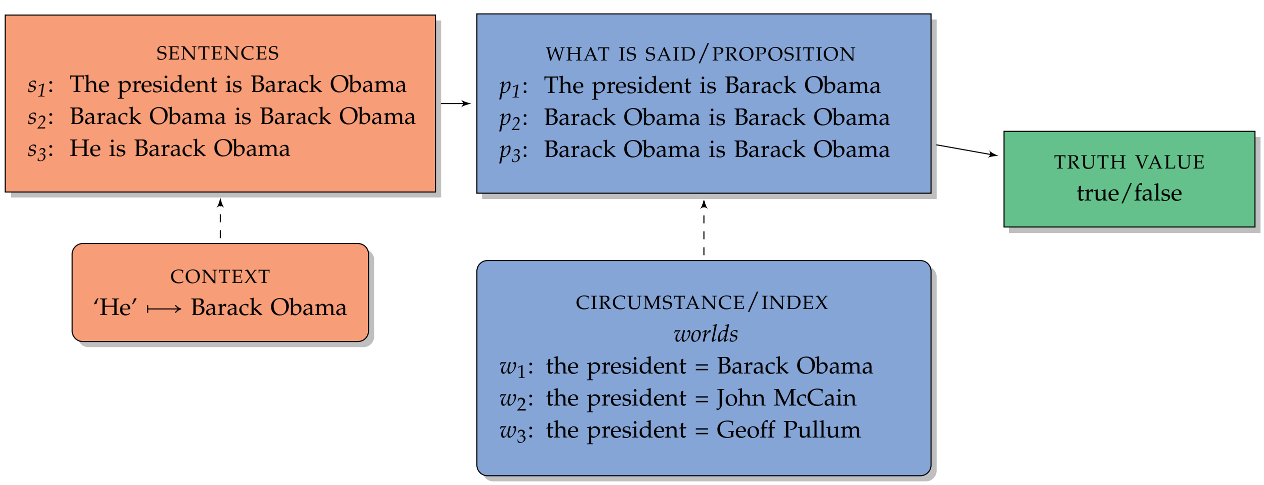

· Look at the descriptive meaning together with a world (CIRCUMSTANCE), and find the individual in that world that satisfies the description.

  · For example, "The president of the United States" denotes Barack Obama at the actual world.

- A definite description is a complex expression whose semantic value depends on the CIRCUMSTANCE (i.e. the possible world at which the sentence is being evaluated for truth).

9.3.3 The Meaning of Indexicals

- Indexicals do seem to have some kind of descriptive meaning — for example:

  'I' = the person speaking.

  'you' = the person addressed by the speaker.

  'now' = the time of the context.

- So, why not just think that indexicals are really just covert definite descriptions?
- Let's test that idea. Suppose we assume that the meaning of the indexical 'I' is as given in (239).

(239) 'I' is shorthand for 'the person who is speaking'. So:
$[\![\text{I}]\!]^{c,w} = \iota x.\text{person-speaking}(x)$

- With this analysis in hand, we might then think that to evaluate any sentence containing the indexical 'I', we should look at each world to see which individual is speaking at that world (just like we would with any other definite description).
- Here's an example: Suppose Caroline says (240).

(240) I am a lecturer.

- Question: Assume that everything in the world stays fixed except for the assumption that Caroline actually uttered (240). Given this, is what is said by (240), had she uttered it, true or false?
  Intuitive answer: True. In a world where she does not utter this sentence but everything else stays fixed, the proposition expressed by 'I am a lecturer' as asserted by Caroline is true, because she is a lecturer in that world.

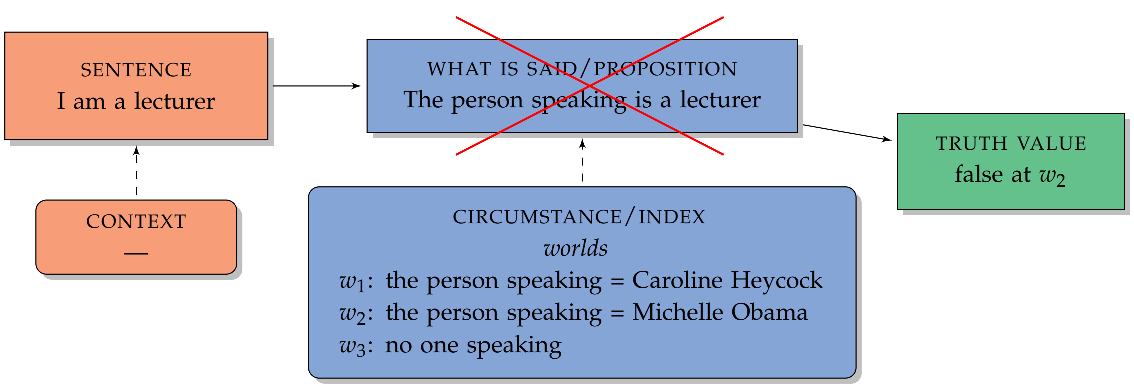

9.4 Adding and Shifting Parameters of the Circumstance

- It seems natural to take sentences to be not just about the actual world, but also about the time of the context and the location of the context. Consider e.g. (241)

(241) It is snowing.

- If I assert (241), it seems that I am making an assertion about what the actual world is relative to the time of my utterance and relative to the location Edinburgh.
- However one could sensibly ask about whether what I asserted would be true in some counterfactual circumstance, e.g.

(241) a. Would (241) be true if more CO_2 had been emitted?
b. Would (241) be true in two days?
c. Would (241) be true in Paris?

- Of course, these kinds of questions might be rare in actual everyday communication.
- However, consider the following:

EVALUATING SENTENCES FOR TRUTH
To determine the truth value of a sentence (asserted in some context), we need to know the relevant evaluation points, namely at which **world**, **time**, and **location** to evaluate it.

For example, (1*) might be true at the world of the context c (i.e. the actual world), the location of c, and the time of c, but false at the world of c, the location of c, and **some past time**.

(1*) The students are ready to learn about *Demonstratives*.

Normally (absent any temporal, locative, or modal shifting, cf. above) the relevant evaluation points are: the actual world (@), the time of the context, and the location of the context.

Given this, we should relativize interpretation to more parameters, i.e. where ϕ is a sentence, then instead of just writing $[\![\phi]\!]^{c,w}$, we should write something like: $[\![\phi]\!]^{c,\langle w,t,l\rangle}$

The function of temporal, locative, and modal expressions/operators is to shift the points at which we evaluate the embedded sentences for truth. See examples below.

- Now, consider the following pairs of sentences:

(242) David is teaching *Demonstratives*.
(243) David **will be** teaching *Demonstratives*. (temporal shifting)

(244) It is raining.
(245) It is raining **in Paris**. (locative shifting)

(246) Jack is the killer.
(247) **It is possible** that Jack is the killer (modal shifting)

- So, how are the truth conditions of a sentence such as e.g. (243) affected by the presence of the temporal operator? Very roughly as follows:
  - "David will be teaching *Demonstratives*" is true at a context c iff "David is teaching *Demonstratives*" is true world of c, the location of c, and at *some time* t' (where t' is later than the time of c).
- Or, more succinctly:
  - $[\![\text{Will}(\phi)]\!]^{c,\langle w,t,l\rangle} = 1$ iff $\exists t'$ such that $t' > t$ and $[\![\phi]\!]^{c,g,\langle w,t',l\rangle} = 1$

Assume that '$x > y$' means that y temporally precedes x. That is, x occurs later than y.

- What about (247)? Well, again very roughly:
  - "It is possible that Jack is the killer' is true at a context c iff "Jack is the killer" is true at the location of c and the time of c at *some* possible world.
- Or, more succinctly:
  - $[\![\text{Possibly}(\phi)]\!]^{c,\langle w,t,l\rangle} = 1$ iff $\exists w'$ such that $\mathcal{R}(w,w')$ and $[\![\phi]\!]^{c,g,\langle w',t,l\rangle} = 1$

9.4.1 Indexicals, Context and Shifting.

- One of Kaplan's key insights is that the semantic values of indexicals and demonstratives is always determined by context and moreover *in*sensitive to these "shifty" operators. For example:
  - **Temporal Indexicals**
    Compare the following two sentences.

    (248) In 100 years, most politicians alive are corrupt.
    (249) In 100 years, most politicians alive now are corrupt.
  - **Locative Indexicals**
    And compare the following two sentences.

    (250) In France, everyone loves the food.
    (251) In France, everyone loves the food here.
- So, we need a way of capturing the following things:

· To evaluate a sentence for truth/falsity it must be evaluated relative to a time, location, and world (which might have been shifted if the sentence is embedded under an intensional operator).

· The semantic value of indexicals remain fixed to the time, the location, and the world of the context of utterance.

9.5 Character and Content

- Kaplan distinguishes two different types of meaning, CHARACTER and CONTENT.

9.5.1 Content

So, the terms CONTENT, WHAT-IS-SAID, PROPOSITION, and INTENSION are being used roughly interchangeably here. There are subtle differences between these terms though.

- On Kaplan's view, the CONTENT of a sentence is the literal information conveyed by the sentence after the meaning of indexicals/demonstrative expressions have been resolved. Following Grice (1989), Kaplan also refers to this level of content as *what is said*.
- In other words, one computes the CONTENT of a sentence by pairing it with an appropriate context of utterance.

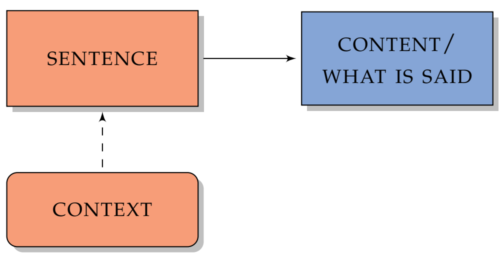

9.5.2 Contents as Functions

The EXTENSION of an expression is its denotation at the world of evaluation (i.e. the actual world). The INTENSION is a function from worlds to EXTENSIONS

- Relative to sentences, CONTENTS can be thought of as functions from CIRCUMSTANCES-of evaluation (ordered pairs of worlds, times, and perhaps locations) to TRUTH VALUES.
- These types of functions are now more commonly referred to as PROPOSITIONS, but also known as INTENSIONS (contrasting the term EXTENSIONS).

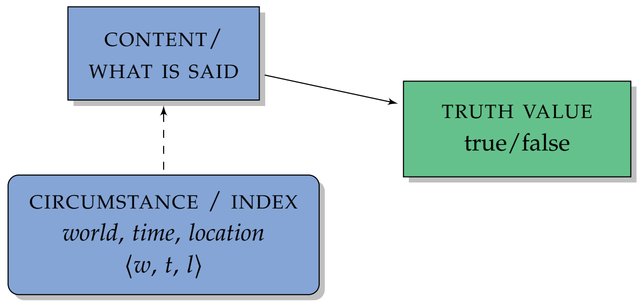

- In short, CONTENT is a function from CIRCUMSTANCE to TRUTH VALUE.

$$\begin{array}{ccc} \textbf{CONTENT}\text{: CIRCUMSTANCE} & \longmapsto & \text{TRUTH VALUES} \\ \uparrow & & \text{(extension)} \\ \langle w, t, l \rangle & & \end{array}$$

9.5.3 Character

- So what is CHARACTER? One can think of the CHARACTER of an expression as a rule that states what content an expression (or sentence) can be used to express relative to different contexts.
- For example, these rules might look something like the following for 'I', 'you', 'she':
  - 'I' refers to the speaker or writer in the context c.
  - 'You' refers to the addressee in the context c.
  - 'Now' refers to the time of the context c.
- The character is (in a sense) the (probably tacit) semantic knowledge that a speaker of the language would need to have in order to be competent with the words.
- **Representing Characters**

  > Just as it was convenient to represent contents by functions from possible circumstances to extensions (Carnap's intensions), so it is convenient to represent characters by functions from possible contexts to contents. (Kaplan, 1989, 505)

- To illustrate, think of CONTEXTS as sequences of objects:

	⟨	SPEAKER	ADDRESSEE	TIME	WORLD	...	⟩
c_1		Anders	Caroline	1pm	@	...	
c_2		Sue	Mary	4pm	@	...	
c_3		Sam	Bob	10am	@	...	
c_n		—	—	—	—	...	

- Now consider the sentence below:

(252) *I* am sending the handout to *you now*.

- To resolve the indexicals in (252), we simply have to inspect their characters (the rules for determining referents). For example:
  - If the sentence in (252) is paired with c_1, then 'I' refers to Anders, 'you' refers to Caroline, and 'now' refers to 1pm.
  - If the sentence in (252) is paired with c_2, then 'I' refers to Sue, 'you' refers to Mary, and 'now' refers to 4pm.
- In other words, the CHARACTER of an expression is a function from CONTEXTS to CONTENTS.

CHARACTER: CONTEXTS $\longmapsto$ CONTENTS
$\uparrow$ (intension)
$\langle$speaker, addressee, time, ... $\rangle$

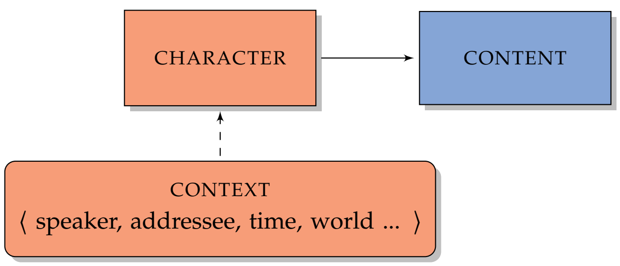

- Summing up, CONTENT and CHARACTER can be thought of in the following simple ways. If ϕ is a sentence:
  - **CHARACTER** of ϕ: $\lambda c.\ \lambda\langle w,t,l\rangle.\ [\![\phi]\!]^{c,\langle w,t,l\rangle}$
  - **CONTENT** of ϕ: $\lambda\langle w,t,l\rangle.\ [\![\phi]\!]^{c,\langle w,t,l\rangle}$

9.6 Stable vs. Unstable (context-sensitive) Characters

- Indexicals have a context-sensitive CHARACTER. That is, the content (semantic value) of an indexical depends on features of the context.
- In contrast, non-indexical expressions (for example 'green', 'car', 'Anders') have stable characters — these words have the same content in every context.
- For example, the sentence in (253) contains no context-sensitive expressions (or so Kaplan claims), so the content of (253) does not vary with contexts. It has the same meaning in all contexts.

(253) It's raining.

- The sentence in (253) is neutral with respect to speaker, addressee, time, location, world, etc.

- This means that if we want to evaluate the content expressed by (253) for truth, the relevant time, location, and world are simply the time, location, and world of the CONTEXT.
- So, if I were to assert (253) now, my assertion would be true iff it is raining at the time, location, and world where I asserted the sentence.
- But, as mentioned earlier, one can shift the parameters of the CIRCUMSTANCE using various operators. For example, whether (254) is true depends on the weather in Paris, not the weather at the location of the context.

(254) In Paris, it's raining.

9.7 Logical and Necessary Truths

- In classical logic, all logical truths are necessary truths.
- For example, a sentence ϕ is a logical truth (a tautology) iff ϕ is true in every model. And in intensional/modal logic, ϕ is a logical truth in system Δ iff ϕ is true in every world of every model of Δ.
- The sentence in (255) has a tautologous feel to it, because it seems that it could never be falsely asserted.
- However, even though (255) will always be true when asserted, it is intuitively clear (255) is not a necessary truth — i.e. (256) is intuitively false.

(255) I am here now.
(256) It is necessary that I am here now. [i.e. things could not have been otherwise]

- Using Kaplan's analysis, we can explain the intuition that (255) is a kind of logical truth (could never be falsely asserted) while avoiding the conclusion that it is necessarily true (i.e. that (256) is true).
- Remember, to evaluate a sentence for truth, any indexicals must first be resolved (using their characters).
    - If the time of the context is t_1, the location of the context is l_1, and the world of the context is w_1, then if the sentence contains no any index-shifting operators, we evaluate the resulting content for truth at t_1, l_1, and w_1.
- So, for example, if I were to assert (255) now, we would get the following result.

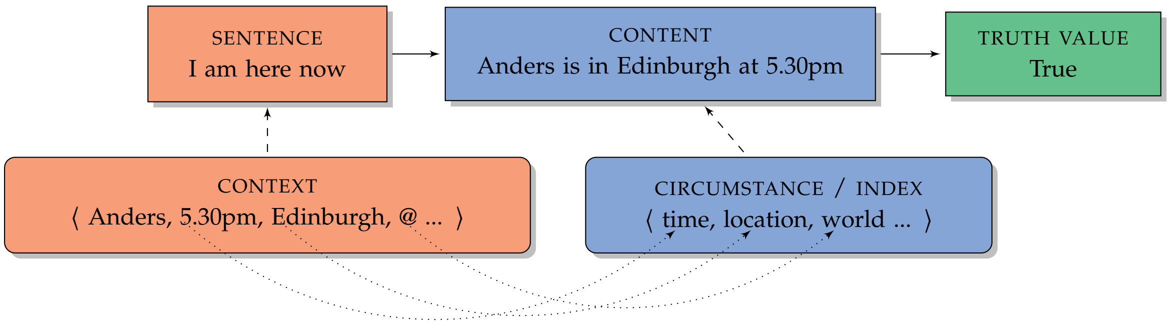

- To obtain the result that (255) is in fact a logical truth, Kaplan redefines truth and logical truth roughly as follows:
  - **Truth**
    If c is a context, an occurrence of a sentence ϕ in c is true iff the content expressed by ϕ in c is true when evaluated with respect to the circumstance of the context.
  - **Logical Truth**
    a sentence ϕ is a logical truth iff ϕ is true relative to every context.
- However, here is why (255) is not a *necessary truth* on Kaplan's analysis, i.e. why (256) will not come out true on Kaplan's analysis. Remember, a claim of the form 'It's necessary that ϕ' is true only if ϕ is true at every possible world.

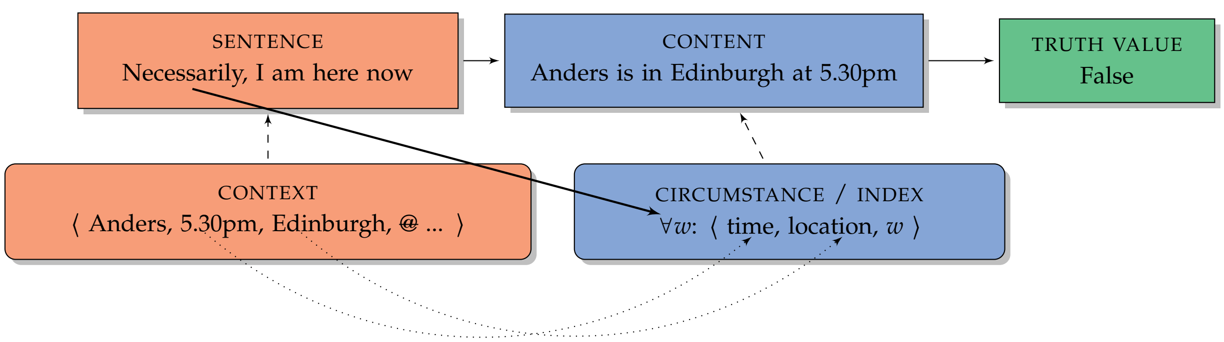

9.8 Monsters

- On Kaplan's semantic framework, to determine the truth value of a sentence, the sentence must be paired with two separate "indices".
  - CONTEXT: ordered pair consisting of speaker, addressee, world, time, location, etc.

$$\text{Formally: CONTEXT} = \langle a, b, w, t, l, \ldots \rangle$$

For this reason, Kaplan's theory is generally referred to as having *double indexing*. However, that double indexing is needed to adequately capture the meaning of indexicals was already pointed out by Kamp (1981).

- CIRCUMSTANCE: an ordered pair consisting of a world, time, location.

  Formally: CIRCUMSTANCE = $\langle w, t, l \rangle$

- **Intensional Operators**
  Furthermore, according to Kaplan, intensional operators, e.g. modals, temporal, and locative expressions, only shift parameters in the CIRCUMSTANCE/INDEX.
- So, one might wonder whether there are any expressions in natural language that have the explicit function of shifting parameters of the CONTEXT?
- Kaplan refers to such expressions, i.e. expressions that shift parameters of the CONTEXT as *monsters*. Moreover, Kaplan explicitly claims that there are no monsters in English and that none could be added either!
- As evidence for Kaplan's claim, consider the following attempts to produce a so-called monstrous operator.

(257) It's true in some context that I'm tired now.
(258) For some speaker in some context, I'm tired.
(259) If Jack is the speaker in this context, then I'm tired.

- Notice that the reference of *I* is not shifted by these operators.

Chapter 10

A Uniform Analysis of Pronouns

- Kaplan's analysis has many virtues, but one immediate problem with this analysis is that it seems to be silent on the issue of bound occurrences of pronouns, e.g. examples such as the following:

(260) [Every girl]$_6$ loves her$_6$ mother.
(261) [Every boy]$_6$ said that he$_6$ passed the exam.

- On Kaplan's analysis, one seems forced to accept that the pronouns in the sentences above are fundamentally different from the pronouns 'her' and 'he' when these are used demonstratively. But this seems very counterintuitive. Evidently these are the same words, they just have different uses.
- To rectify this, we will look at Heim and Kratzer's alternative analysis of pronouns.

10.1 Anaphoric vs. Deictic (uses of) Pronouns

- Each of the sentences below have an indefinite number of potential meanings (due to the presence of a pronoun).

(262) John likes running and he likes skiing.
(263) Mary and John came to the party and they brought wine.
(264) If Susan interviewed, then she will get the job.

- However, we can distinguish between two types of interpretations of pronouns, namely when:

    a. The semantic value of the pronoun is determined by a linguistic antecedent. (anaphoric use)
    b. The semantic value of the pronoun is determined by a demonstration by the speaker. (deictic/demonstrative use)

- There are also so-called *cataphoric* uses of pronouns (cases of anaphoric pronouns where the pronoun precedes its linguistic antecedent). These are quite rare, but here is an example.

(265) As he climbed the mountain, John realized the battle was lost.

- When the semantic value of a pronoun is anaphorically determined, we indicate this by co-indexing. I.e.

(266) John$_1$ likes running and *he*$_1$ likes skiing.
(267) John$_1$ likes running and *he*$_2$ likes skiing.

- The distinction between anaphoric and deictic/demonstrative is not exhaustive, i.e. there are uses of pronouns that do not seem to fit these labels. We will mostly ignore these, but here are some examples.

  · **Pure Indexicals**
  Some pronouns (what Kaplan (1989) labeled 'pure indexicals') do not require linguistic antecedents or demonstrations, e.g.

  (268) *I* am tired.

  · **Contextual Salience**
  Cases where salience in context seems to determine the semantic value of the pronoun. For example, consider this type example due to, I believe, Irene Heim: Imagine you are looking at a line up of soldiers. Suddenly one of the soldiers drops his rifle. Suppose I now utter (269).

  (269) He's going to be in trouble now.

  · **Pronouns of Lazyness**
  Cases where a pronoun seems to go proxy for a full descriptive phrase.

  (270) This year the president is a Republican, but one fine day, he ('the president') will be a member of the Green party.

  · **Paycheck Pronouns** (bound pronouns of lazyness)
  Cases where the pronoun seems not only to go proxy for a full descriptive phrase but where the pronoun also appears to be bound.

  (271) Mary, who deposited her paycheck at the ATM, was smarter than any woman who kept it ('her paycheck') in her purse.

- However, with regards to semantics, the distinction between anaphoric and non-anaphoric uses might not be particularly useful.
- The reason is that anaphoric and deictic pronouns are often treated the same way semantically. A more useful distinction is that between *referential* and *bound* pronouns.

10.1.1 Referential Pronouns

- In a standard case of a deictic pronoun, the pronoun is treated semantically as a referential term. I.e.

(272) He is tired.

- If a speaker utters (272) while pointing to Bob, the semantic value of the pronoun is Bob. In contrast, if the speaker is pointing to Jack, the semantic value of the pronoun is Jack. Hence, the pronoun is *referential* — it refers to a particular individual.
- The standard treatment of referential uses of pronouns is simply that the semantic value of the pronoun is determined by a variable assignment. Hence...

(273) $[\![he_1]\!]^g = g(1)$
(274) $[\![he_4]\!]^g = g(4)$
(275) $[\![he_{23}]\!]^g = g(23)$

- And given this variable assignment...

$$g = \begin{bmatrix} 1 & \longmapsto & \text{Jack} \\ 2 & \longmapsto & \text{Bob} \\ 3 & \longmapsto & \text{Frank} \\ 4 & \longmapsto & \text{Peter} \\ \vdots & \longmapsto & \vdots \\ 23 & \longmapsto & \text{Frank} \\ \vdots & & \vdots \end{bmatrix}$$

- ... we get the following:

(276) $[\![he_1]\!]^g = g(1) = \text{Jack}$
(277) $[\![he_4]\!]^g = g(4) = \text{Peter}$
(278) $[\![he_{23}]\!]^g = g(23) = \text{Frank}$

- The important point here is that some anaphoric uses of pronouns (e.g. those mentioned above) can also be treated as referential.
- On the anaphoric reading of the pronouns in (262)–(264), these pronouns simply refer to whoever their antecedents refer to. For example.
  - In (262), the linguistic antecedent for 'he' is the name 'John'. Since the name 'John' refers to John, the pronoun 'he' *refers* to John too.
  - In (263), the linguistic antecedent for the pronoun 'they' is 'Mary and John'. Since 'Mary and John' can be taken to refer to a group of two people (or a *plural individual)*, the pronoun 'they' can be taken to *refer* as well.
- **In conclusion**, when the antecedent of an anaphoric pronoun is an individual and when the pronoun is used deictically, the pronoun is referential – and thus treated formally as a variable whose semantically value is a function of the variable assignment. That is, whatever the assignment maps the variable to, that is its semantic value.

Interlude: Context of Utterance

- One pertinent question is where exactly variable assignments come from?
- Generally speaking, the variable assignment is a way of modeling the intentions of the speaker.
- On the Kaplanian picture, a context contains, at least, a speaker, an addressee, a world, a time, a location, etc. It seems natural to think that such a context, given that it contains the speaker, also determines the variable assignment.
- So, we assume that context determines a partial function from indices to individuals, i.e. a variable assignment g — and thus that the following appropriateness condition obtains for contexts.

  — **Appropriateness Condition**
  A context c is appropriate for an LF ϕ only if c determines a variable assignment g_c whose domain includes every index which has a free occurrence in ϕ.

- So, our interpretation function will now be relativized to *two* parameters.

$$[\![\cdot]\!]^{c,g}$$

- However, you might be wondering, if the variable assignment is used to determine the meaning of pronouns, referential *and* bound, then why do we also need c? We will answer this question in the next section.

10.2 Pronominal Constraints

- There are many different pronouns in English, e.g. *he*, *she*, *I*, *you*, *they*, but these pronouns seem to have different meanings (different 'stable' meanings to use Kaplan's terminology). More precisely, pronouns seem to have differing constraints on their use, for example...

$$\text{`he' must refer to an individual who is} \begin{cases} \text{MALE} \\ \text{SINGULAR} \\ \text{NOT SPEAKER NOR ADDRESSEE IN } c \end{cases}$$

$$\text{`she' must refer to an individual who is} \begin{cases} \text{FEMALE} \\ \text{SINGULAR} \\ \text{NOT SPEAKER NOR ADDRESSEE IN } c \end{cases}$$

$$\text{`I' must refer to an individual who is} \begin{cases} \text{SINGULAR} \\ \text{THE SPEAKER IN } c \end{cases}$$

10.2.1 Characters vs. ϕ-Features

- Kaplan's way of capturing this 'stable' part of the meaning of indexicals (e.g. pronouns) was by appealing to characters.
- Remember, a character is a function from a context to a content — for example, the character of e.g. 'I' is a function from a context to the speaker in that context.
- Also, according to Kaplan, characters *determine* reference (relative to a context c). In other words, characters are assumed to provide necessary *and* sufficient conditions for reference.
- **Note**: While it is (perhaps) easy to explicate the character of 'I' (which *determines* its referent), it's much harder to explicate it for many other pronouns, e.g. 'he'. Remember, the function (relative to a context) must *determine* the referent.
- We are not going to use Kaplan's notion of character here, however we do want to capture the stable meaning of various pronouns. To do this, we exploit the fact that pronouns are associated with grammatical features, so-called ϕ-features (person, number, gender features, cf. above — also called *PNG*-features).
- ϕ-features are generally assumed to be syntactically represented, i.e. as follows.

(279)

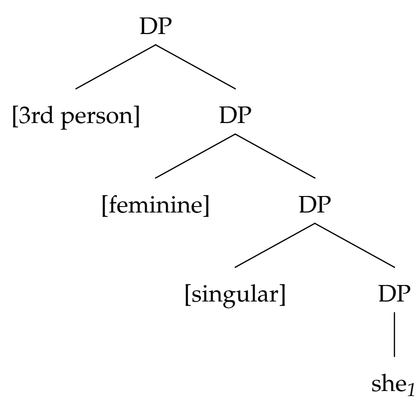

- The lowest DP-node is interpreted using the Pronouns and Traces rules, and the higher nodes by FA.
- The features are then given the following lexical entries:
  - $[\![\text{feminine}]\!] = [\lambda x_e\text{: } x \text{ is female . } x]$
  - $[\![\text{masculin}]\!] = [\lambda x_e\text{: } x \text{ is male . } x]$
  - $[\![\text{1st person}]\!] = [\lambda x_e\text{: } x \text{ is the speaker in } c \text{ . } x]$
  - $[\![\text{2nd person}]\!] = [\lambda x_e\text{: } x \text{ is the addressee in } c \text{ . } x]$
  - $[\![\text{3rd person}]\!] = [\lambda x_e\text{: } x \text{ is not the speaker in } c \text{ nor the addressee in } c \text{ . } x]$
  - $[\![\text{singular}]\!] = [\lambda x_e\text{: } x \text{ is an atomic individual . } x]$
  - $[\![\text{plural}]\!] = [\lambda x_e\text{: } x \text{ is a plural individual . } x]$

- These are partial identity functions — each function places a constraint on the individual a assigned to the variable by the variable assignment and returns a if the constraint is satisfied.
- On this analysis, ϕ-features are essentially treated as if they were *presuppositions*.
- As regards the question, why do we need a parameter for the context c? Notice that the variable assignment is assumed to be determined by c and, more importantly, that some of the ϕ-features make explicit reference to c. This is why a parameter for the context c is needed **in addition** to the variable assignment g.

Remember, an identity function, e.g. $[\lambda x_e.\ x]$ is a function that takes an argument of some type τ and then outputs that very same argument. So, if the input to $[\lambda x_e.\ x]$ is Jack, the output is just Jack, viz. $[\lambda x_e.\ x]$(Jack) = Jack.

10.2.2 Bound Pronouns

- We have already encountered several cases where a pronoun is bound by a higher quantifier and this will generally be the case every time a pronoun is co-indexed with a quantificational noun phrase. For example,

(280) [No boy]$_1$ lost his$_1$ wallet.
(281) [Every boy]$_1$ saw himself$_1$.
(282) [Some boy]$_1$ thought he$_1$ could walk on water.

- A movement analysis of these cases were discussed in earlier notes.
- If the pronouns in the above sentences were not co-indexed with the quantifier (except for the reflexive pronoun in (281) which must be co-indexed), these pronouns would have to be interpreted as deictic pronouns and thus have to be interpreted relative to a particular assignment.

Bound or Referential Analysis?

- When a pronoun is co-indexed with a non-quantificational antecedent, it would seem that whether it is bound or referential makes no truth conditional difference.

(283) Bob loves his mother.

(284)

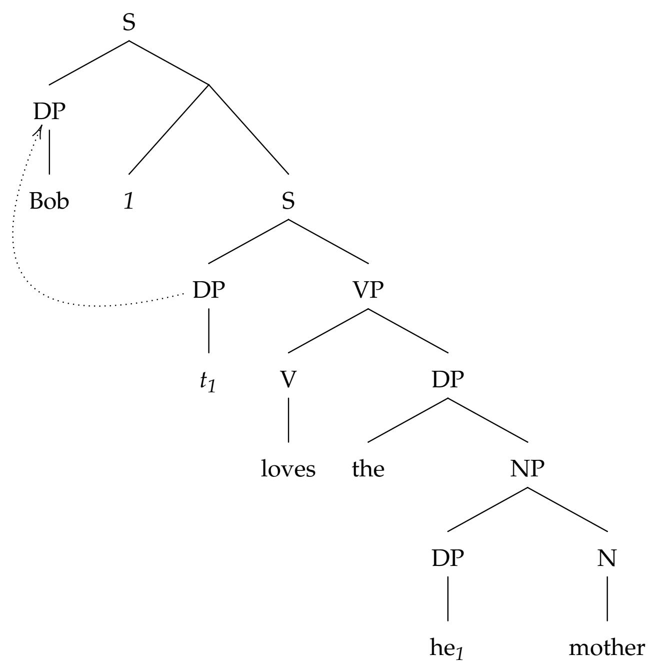

Notice that for the purpose of the following discussion, we need to assume that "mother" is a so-called relational noun, namely [λy_e . λx_e. x is a mother of y]

- In (284), the DP has been QR'ed and adjoined to S. So, a trace co-indexed with the binder is left behind.
- This means that in (284), the pronoun ends up getting bound, and the embedded S ends up denoting the following property: [λx_e . x loves the mother of x].

It seems that we need to assume that QR'ing the DP has the effect of making the DP control what the variable assignment assigns to the index that drops out — otherwise we would get presupposition failure for 'his mother'. However, even when making that assumption, it's not clear how to deal with sentences such as "Every boy loves his mother".

- In contrast...

(285)

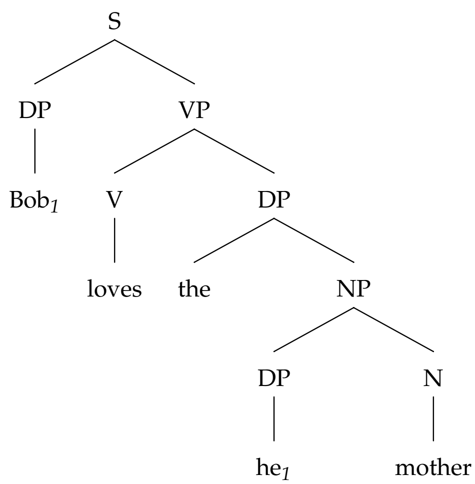

- Here the pronoun must be interpreted via a variable assignment.
- But since it is co-indexed with Bob, g maps the pronoun to Bob — the VP thus ends up denoting the following property: [λx_e . x loves the mother of Bob]
- It seems that with our current rules, we end up predicting a great deal of ambiguities that have zero truth conditional effects. This does not look like a theoretically parsimonious result.
- However, there are good reasons to think that we actually need both these LFs.

10.2.3 VP-Ellipsis: Strict and Sloppy.

- The following sentence has (at least) two interpretations.

(286) John loves his mother and Bill does too.
  a. John loves John's mother and Bill loves John's mother too (strict)
  b. John loves John's mother and Bill loves Bill's mother too (sloppy)

- Before explaining what is going on in these cases, let's quickly consider what is called the inverted Y model in Government and Binding (G&B) theory.

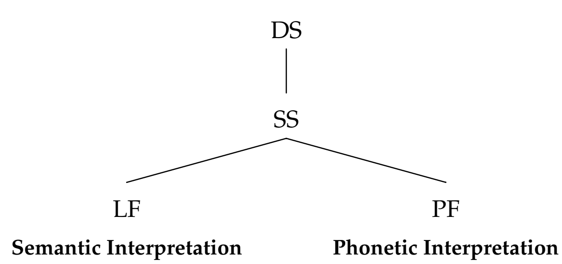

- We are assuming that the input for semantics is LFs. In G&B theory, LFs are derived via transformational rules from surface structures (SS). In contrast, PF refers to the phonetic representation of the sentence to the speaker, i.e. the audible (or in the case of sign language, a visible) representation of the sentence. The input for the derivation of PF is also surface structure.
- We do not need to dwell on the intricacies of LF versus PF, but it is important to realize that the LF and PF associated with a sentence may differ substantially.

Constraints on Ellipsis

- In cases of VP-ellipsis, it is generally assumed that the elided material is only elided at PF, i.e. the elided material is simply not pronounced.
- However, this elided material is fully represented at LF.
- So, what occurs at LF (but is elided at PF) is an exact and fully disambiguated copy of the material in the preceding sentence. To appreciate why, consider the following sentences.

(287) Laura showed a drawing to every teacher.

a. $\exists x[\text{drawing}(x) \wedge \forall y[\text{teacher}(y) \rightarrow \text{showed}(\text{Laura}, x, y)]]$
b. $\forall y[\text{teacher}(y) \rightarrow \exists x[\text{drawing}(x) \wedge \text{showed}(\text{Laura}, x, y)]]$

(288) Laura showed a drawing to every teacher but Lena didn't.

- Depending on the interpretation of the first conjunct, the second conjunct (with material elided) must be interpreted in the same way.

LF Identity Condition on Ellipsis

- A constituent may be deleted at PF only if it is a copy of another constituent at LF.
- Now consider again (286).

(286) John loves his mother and Bill does too.

- In (286), the constituent which is the sister of the 'DP' on the left is elided. Now consider what happens if we copy over this constituent, cf. (285) below.

(285)

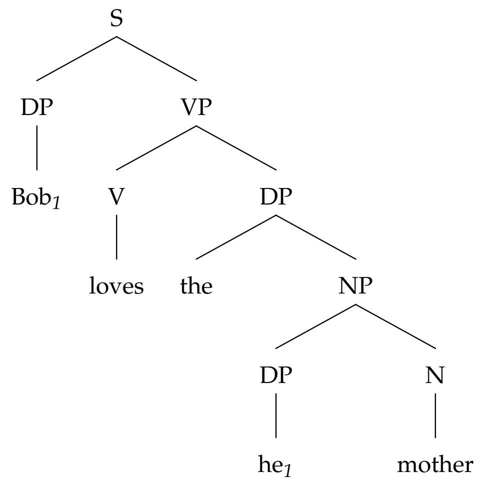

- Remember, this constituent denotes the following property: [λx_e . x loves Bob's mother]
- In contrast, consider what happens if we copy over the elided constituent of (284).

(284)

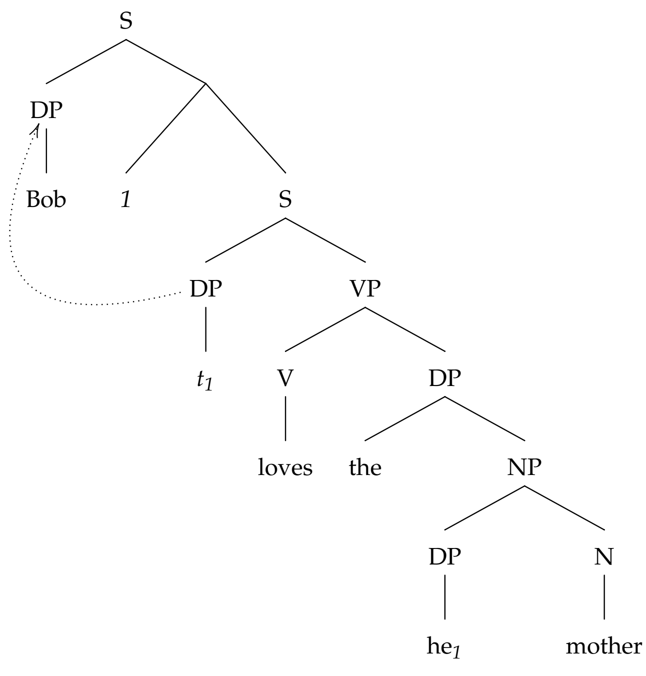

- Remember, this VP is going to end up denoting the following property: [λx_e . x loves x's mother]
- Hence, we are going to need both the representations in (284) and (285) to capture both sloppy and strict readings of sentences such as (286).

Chapter 11

Proper Names Revisited

In progress

11.1 Millianism

- Kripke's Semantic and Modal Arguments
- Rigidity
- Direct Reference

11.1.1 Three Challenges for Millianism

11.2 Predicativism

- *That*-Predicativism
- *The*-Predicativism

11.3 Variabilism

- Names as indexicals
- Names as pronouns

Chapter 12

Intensional Semantics

12.1 Beginnings

12.1.1 Displacement

- As nicely demonstrated by Kaplan, natural language is not restricted to discourse about the *actual here* and *now*. This feature of natural language is sometimes referred to as *displacement*.
- Examples include the familiar cases of locative and temporal displacement.

(289) It is raining in Edinburgh.
(290) In two weeks, it will be snowing.

- However, locative and temporal displacements are just some of a wide variety of types of displacement. For example, natural language also permits talk about e.g. *counterfactual*, *necessary*, *possible*, or *probable* states of affairs. In addition, we can talk about how others take (or desire) the world to be, and about how the world *generally* is — just to mention a few.

(291) **Counterfactuals**
If hurricane Sandy had gone straight north, it would have rained in Boston.

(292) **Modal Adverbs**
It's necessary/possible that it's raining in Edinburgh.

(293) **Modal Auxiliaries**
It might/must/should/ought to be raining in Edinburgh
It is probably raining in Edinburgh.

(294) **Propositional Attitudes**
Jaakko believes/hopes/dreams that it's raining in Edinburgh.

(295) **Habituals**
Jaakko smokes.

(296) **Generics**
Birds fly.

- Intensional semantics is designed to capture displacement, i.e. to capture the meaning of temporal, modal, and other displacing operators.

12.2 Extending the Semantic System to Modals

- To extend the analysis pursued in chapters 1-9 to e.g. modals, we will make a simplifying assumption, namely that both modal auxiliaries and modal adverbs are sentential (or propositional) operators.
- I.e. we assume that the logical form of sentences like (297) is (297a) — or more conspicuously (297b). In other words, we will assume that modal expressions are always raised and adjoined to S.

(297) It must be raining in Copenhagen.
a. must(it is raining in Copenhagen)
b. must(ϕ)

(298) It might be raining in Copenhagen.
a. might(it is raining in Copenhagen).
b. might(ϕ).

- Some central questions to be addressed then include:
  - What function should we assume that these expressions denote?
  - What should we assume about the arguments of these functions?

12.2.1 Non-Truth Functional Operators

- Since expressions such as 'it is necessary that' and 'it is possible that' appear to combine syntactically with complete sentences (*that*-clauses), it might seem naturally to think that these should be treated as unary truth functional operators (like e.g. negation).
- However, it is very easy to see that no truth functional analysis can correctly capture the meaning of e.g. 'might', 'possibly' or 'it is possible that'. To demonstrate, consider what function 'might' would have to denote if it was a truth functional operator.
- First, intuitively, if a sentence ϕ is true, it seems that 'might ϕ' should also be true. So, 'might' should map true sentences to true.

ϕ	T	F
might(ϕ)	T	

- The question is what the function should output when the input is false. There are two options:

  (a) **Mapping 'might ϕ' to *false* if 'ϕ' is false.**
  Suppose you roll a fair die. Before looking to see what number you rolled, you now say "it might be a six". This seems true. It seems true even if you later discover that it is not a six.

  In short, we will get incorrect results whenever the embedded sentence is false, but contingent (i.e. not necessarily false).

  (b) **Mapping 'might ϕ' to *true* if 'ϕ' is false.**
  Suppose that ϕ is "some bachelors are married". That sentence is obviously false, but so, intuitively is "some bachelors might be married".

  In short, we will get incorrect results whenever the embedded sentence is necessarily false.

- In conclusion, modal expressions cannot simply be treated as truth functional operators, but since their arguments appear to be sentences (or propositions), we seem to need to move beyond standard extensional semantics and add an extra layer of complexity to our semantic framework.

12.3 Possible Worlds

- When discussing Kaplan's framework, we introduced possible worlds into the semantics. However, we did not discuss what exactly possible worlds are.
- One way to think about possible worlds is simply as a way that the actual world *could*, *should*, or *would have to be* like.
- Here is a succinct quote from Lewis (also cited in von Fintel and Heim):

> The world we live in is a very inclusive thing. Every stick and every stone you have ever seen is part of it. And so are you and I. And so are the planet Earth, the solar system, the entire Milky Way, the remote galaxies we see through telescopes, and (if there are such things) all the bits of empty space between the stars and galaxies. There is nothing so far away from us as not to be part of our world. Anything at any distance at all is to be included. Likewise the world is inclusive in time. No long-gone ancient Romans, no long-gone pterodactyls, no long-gone primordial clouds of plasma are too far in the past, nor are the dead dark stars too far in the future, to be part of the same world. [...]
>
> The way things are, at its most inclusive, means the way the entire world is. But things might have been different, in ever so many ways. This book of mine might have been finished on schedule. Or, had I not been such a commonsensical chap, I might be defending not only a plurality of possible worlds, but also a plurality of impossible worlds, whereof you speak truly by contradicting yourself. Or I might not have existed at all — neither myself, nor any counterparts of me. Or there might never have been any people. Or the physical constants might have had somewhat different values, incompatible with the emergence of life. Or there might have been altogether different laws of nature; and instead of electrons and quarks, there might have been alien particles, without charge or mass or spin but with alien physical properties that nothing in this world shares. There are ever so many ways that a world might be: and one of these many ways is the way that this world is. (Lewis, 1986, 1)

- With respect to the semantic framework outlined in chapters 1-9, we now introduce a set of possible worlds, namely the set of *all* possible worlds. We will refer to this set as $\mathcal{W}$. An interpretation of an expression or full sentence will therefore be relativized to a possible world.

12.4 A Toy Example: Sherlock Holmes

- Consider the sentence below.

(299) A famous detective lives at 221B Baker Street.

- Relative to a world and a variable assignment, we get these truth conditions.

(300) $[\![$A famous detective lives at 221B Baker Street$]\!]^{g,w}$ = 1 iff a famous detective lives at 221B Baker Street in world w.

- Here w serves as the *evaluation world* — the world at which we are evaluating the sentence for truth.
- In general, when there are no modal or temporal operators around, the world of evaluation is the world determined by the context, viz. the actual world.

We will ignore context-sensitivity and therefore suppress the c parameter. However, we assume that it is the context that determines the world of evaluation.

12.4.1 World-dependent and world-independent lexical entries

- Just as was the case when we introduced variable assignments, not all lexical entries have world-dependent semantic values.
- For example, we will assume that e.g. proper names, logical connectives, and quantificational determiners have world-*independent* semantic values.
  - $[\![\textit{Noam Chomsky}]\!]^{g,w}$ = Noam Chomsky
  - $[\![\textit{Barack Obama}]\!]^{g,w}$ = Barack Obama
  - $[\![\textit{Hillary Clinton}]\!]^{g,w}$ = Hillary Clinton

  - $[\![\textit{and}]\!]^{g,w} = [\lambda \phi_t \,.\, [\lambda \psi_t \,.\, \phi = 1 \text{ and } \psi = 1]]$
  - $[\![\textit{every}]\!]^{g,w} = [\lambda P_{\langle et\rangle} \,.\, \lambda Q_{\langle et\rangle} \,.\, \text{for all } x \in D_e, \text{ if } P(x) = 1 \text{ then } Q(x) = 1]]$
  - $[\![\textit{a}]\!]^{g,w} = [\lambda P_{\langle et\rangle} \,.\, \lambda Q_{\langle et\rangle} \,.\, \text{there is some } x \in D_e \text{ such that } P(x) = 1 \text{ and } Q(x) = 1]]$
- The expressions that *do* have world-dependent semantic values include predicates such as the following.
  - $[\![\textit{famous}]\!]^{g,w} = [\lambda x_e \,.\, x \text{ is famous in } w]$
  - $[\![\textit{detective}]\!]^{g,w} = [\lambda x_e \,.\, x \text{ is a detective in } w]$
  - $[\![\textit{lives-at}]\!]^{g,w} = [\lambda y_e \,.\, [\lambda x_e \,.\, x \text{ lives at } y \text{ in } w]]$

- Since our lexical entries are relativized to worlds, we are forced to amend our interpretation rules. For example, we could amend the rule of functional application as follows:

> **Functional Application (FA)**
> If α is a branching node $\{\beta,\gamma\}$ the set of of its daughters, then, for any world w and assignment g: if $[\![\beta]\!]^{g,w}$ is a function whose domain includes $[\![\gamma]\!]^{g,w}$, then $[\![\alpha]\!]^{g,w} = [\![\beta]\!]^{g,w}([\![\gamma]\!]^{g,w})$

12.5 Intensional Operators

- So far, the introduction of a world parameter makes very little difference to our predictions, since every sentence will just be evaluated at the actual world.
- However, the point of introducing possible worlds is to capture the meaning of various operators whose function, it seems, is to shift this parameter. I.e. force evaluation of the sentence or expression to some world other than the world of the context.
- Consider a modal operator such as: *in the world of Sherlock Homes*. This operator combines syntactically with complete sentences.

(301) In the world of Sherlock Holmes, a famous detective lives at 221B Baker Street.

- Our semantics should output truth conditions that look something like the following:

(302) $[\![$*In the world of Sherlock Homes, a famous detective lives at 211B Baker Street*$]\!]^{g,w} = 1$ iff the world w' as it is described in the Sherlock Holmes stories is such that there exists a famous detective in w' who lives at 221B Baker Street.

- We can make this more general. We can introduce a rule which for any arbitrary sentence ϕ yields (more or less) accurate truth conditions when ϕ is combined with the operator *in the world of Sherlock Holmes*.

(303) $[\![$*In the world of Sherlock Holmes* $\phi]\!]^{w,g} = 1$ iff the world w' as it is described in the Sherlock Holmes stories is such that $[\![\phi]\!]^{w',g} = 1$.

- Ideally, we would also want the operator *in the world of Sherlock Holmes* to combine via functional application with a sentence — but now we run into a problem.
  - Sentences are of type t.
  - An operator such as *in the world of Sherlock Holmes* (which must be a function that takes sentences as inputs) should output something of type t too.
  - It then follows that this operator would be a function of type $\langle t,t\rangle$.
  - But this is just a standard truth function (like e.g. negation) — and we know that this will not work.

12.5.1 Intensions

- The solution to the above problem is use *intensions* rather than extensions. Remember,

  - The extension (or *semantic value*) of a sentence ϕ relative to a world w and variable assignment g is a truth value, viz. $[\![\phi]\!]^{g,w} \in \{0,1\}$

- However, from $[\![\phi]\!]^{g,w}$, the intension is easily derived. Remember, the intension of a sentence ϕ is simply the set of worlds where ϕ is true. Call that set X. Since $X \subseteq \mathcal{W}$, X has a corresponding characteristic function.
- Hence, the intension of a sentence ϕ is the characteristic function of X over $\mathcal{W}$ — viz. a function from worlds to truth values.
- We use λ-notation to indicate that we are talking about *intensions* rather than *extensions*.

  The intension of ϕ: $[\lambda w.\ [\![\phi]\!]^{g,w}]$

- Since some sentences trigger presuppositions, e.g. sentences containing 'the', some sentences will not have truth values at some worlds. This means that intensions are really partial functions.
- So, strictly speaking, we should write:

  $[\lambda w\text{: } \phi \in \textbf{dom}([\![\cdot]\!]^{g,w}).\ [\![\phi]\!]^{g,w}]$

Strictly speaking, we should also indicate what type of expression this function is defined over — we will add that detail in a bit.

12.5.2 Enriching the Type Theory

- We now have what appears to be a new *type* of expression in our language (intensions), so we need to have appropriate type assignments for such expressions.
- Currently, our inventory of denotational domains consists of two domains for basic types, e and t — and a recursively defined domain for complex types, viz.

  a. D_e = the set of all possible individuals (type e expressions)

  b. $D_t = \{0,1\}$, the set of truth values (type t expressions)

  c. If σ and τ are semantic types, then $D_{\langle\sigma,\tau\rangle}$ is the set of all functions from D_σ to D_τ (type $\langle\sigma,\tau\rangle$ expressions).

- We now add another recursively defined domain.

  d. If σ is a type, then $D_{\langle s,\sigma\rangle}$ is the set of all functions from $\mathcal{W}$ to D_σ

- Functions of type $\langle s, \ldots \rangle$ are intensions.
- Given the above recursive rule, we can talk about the intensions of any expression, e.g.

  - **The intension of a sentence is of type $\langle s,t\rangle$.**
    Intensions are often just equated with *propositions*. As mentioned before, note that intensions of sentences are partial functions.

- **The intension of an intransitive verb is of type** $\langle s,\langle e,t\rangle\rangle$
  These are the intensional versions of properties.
- **The intension of a transitive verb is of type** $\langle s,\langle e,\langle e,t\rangle\rangle\rangle$
  These are the intensional versions of 2-place relations.
- **The intension of expressions of type** e **are** $\langle s,e\rangle$.
  These are functions from worlds to individuals — normally called *individual concepts*.

- **Note**: Perhaps surprisingly, there is no basic type s in our inventory of denotations. The reason for this is that there seems to be no expressions of natural language that refer to possible worlds — the notion of a possible world is a technical notion, a term of art, that we are simply using in our metalanguage to describe truth conditions.
- However, one could add a denotational domain of possible worlds to the type theory which would mean that there was also a basic expressions of type s. Such a type theory is called a two-sorted type theory (or **Ty2**).
- Notice also that for any of this to make any sense, we have to expand our model. Previously a model for the language looked as follows.

$$\mathfrak{M} = \langle D_e, D_t, \mathfrak{I}\rangle$$

- But since we are now assuming that there are intensional meanings — which in the case of propositions are functions f from worlds w to truth values $\{0,1\}$, i.e. $f\colon \mathcal{W} \longmapsto D_t$, our model must be amended as follows.

$$\mathfrak{M} = \langle D_e, D_t, \mathcal{W}, \mathfrak{I}\rangle$$

12.5.3 Semantics for a Shifter

- With our newly revised type theory, we can now formulate the lexical entry for the modal operator *in the world of Sherlock Holmes*.

(304) $[\![$*in the world of Sherlock Holmes*$]\!]^{g,w} = \lambda p_{\langle st\rangle}$. the world w' as it is described in the Sherlock Holmes stories is such that $p(w') = 1$

- This operator is now defined as a function that takes a proposition p (the intension of a sentence ϕ) as argument and outputs a truth value.
- However, we need to add a new rule now for *intensional* functional application.

Intensional Functional Application
If α is a branching node and $\{\beta,\gamma\}$ the set of α's daughters, then, for any world w and assignment g: if $[\![\beta]\!]^{g,w}$ is a function whose domain contains $[\lambda w_{\langle s,\sigma\rangle}.\ [\![\gamma]\!]^{g,w}]$, then $[\![\alpha]\!]^{g,w} = [\![\beta]\!]^{g,w}([\lambda w_{\langle s,\sigma\rangle}.\ [\![\gamma]\!]^{g,w}])$

(assume that σ is a variable ranging over both basic and complex types)

- To keep things simple, the only intensions that will figure in our semantic derivations are propositions. For example, intensions of predicates will not play any role (for now).

12.5.4 Anchoring to the Actual World

- It is not clear that it would be interesting in any sense to know what is going on in other possible worlds unless it was somehow anchored to what is actually going on.
- What our current analysis of e.g. 'in the world of Sherlock Holmes' fails to capture is that the truth value of modal sentences depends importantly not only on what is the case at the other possible worlds, but also on what is the case at the *actual* world.
- To see why this is important, notice that the claim that (305) is a contingent truth. Sherlock Holmes could just as well have lived somewhere else. Similarly, (306) is a contingent falsehood.

(305) In the world of Sherlock Holmes, a famous detective lives at 221B Baker Street.
(306) In the world of Sherlock Holmes, a famous detective lives at 221B Princess Street.

- The fix for this problem is to relativize the meaning of modal operators to the evaluation world, i.e.

(307) $[\![\textit{in the world of Sherlock Holmes}]\!]^{g,w} = [\lambda p_{\langle s,t\rangle}$. the world w' as it is described in the Sherlock Holmes stories in w is such that $p(w') = 1]$

- Now sentences of the form "in the world of Sherlock Holmes, ϕ" make a claim about the evaluation world (i.e. the actual world), namely that some possible world that stands in a *certain relation* to the actual world, has such-and-such properties.

12.5.5 More Worlds

- It is implicitly assumed above that we can restrict our semantics to just one world, namely *the* world of Sherlock Holmes. But this assumptions is not feasible. We really need a set of worlds compatible with the Sherlock Holmes stories in the evaluation world.
- So, instead of treating the operator *in the world of Sherlock Holmes* as *shifting* the evaluation world parameter, we should rather think of it as universally quantifying over a certain set of worlds, i.e.

(308) $[\![\textit{in the world of Sherlock Holmes}]\!]^{g,w} = [\lambda p_{\langle s,t\rangle}$. $\forall w'$ such that w' is compatible with the Sherlock Holmes stories in w: $p(w') = 1]$

Chapter 13

Propositional Attitudes

13.1 Propositional Attitude Verbs as Quantifiers over Possible Worlds

- Propositional attitude ascriptions are sentences where a mental state of some kind is ascribed to one or more individuals, cf. below.

(309) Bill believes that Rome is in France.
(310) Mary hopes that France will win the world cup.
(311) Sue dreamt that she won the lottery.

- Propositional attitude verbs ('believe', 'hope', 'want', 'know', 'suppose', 'desire', ...) are standardly assumed to express a relation between an agent and a proposition. For example, in (309), the verb 'believe' expresses a relation between the agent Bill and the proposition that Rome is in France.
- One way of modeling the meaning of expressions such as 'believe' is to treat them along the lines of our treatment of the operator 'in the world of Sherlock Holmes'. For example, for (309) to be true, it does not have to actually be the case that Rome is in France it just has to be the case that *Bill believes* that Rome is in France.
- So, we might say that for a sentence such as (309) to be true, it must be the case that for every world compatible with what Bill believes, Rome is in France.
- Or, alternatively, in the worlds of Bill's beliefs, Rome is in France. In other words, we might propose the following lexical entry for 'believe'.

(312) $[\![\textit{believe}]\!]^{g,w} = [\lambda p_{\langle s,t\rangle}.\ [\lambda x_e.\ \forall w'$ such that w' is compatible with x's beliefs in w: $p(w') = 1]]$

- This lexical entry assumes that...

  (a) We can think of an agent's beliefs as a set of worlds.
  (b) We can provide a suitable explication of the notion of a world being 'compatible' with a belief.

13.1.1 Beliefs as Sets of Worlds and Compatibility

- Suppose some agent A believes only the following two propositions.
  - Bill snores.
  - Germany won the world cup.
- The set of worlds X_A that characterize A's beliefs is then the following:

$$X_A = \{w \mid \text{Bill snores in } w \text{ and Germany won the world cup in } w\}$$

- Of course, notice that any *logical consequence* of these propositions are now also predicted to be believed by A. For example:
  - For every world w such that $w \in X_A$: there is *someone* who snores. So the proposition that someone snores must also be believed by A.
  - Since logical and mathematical truths are true at every world, it is predicted that A believes *every* logical and mathematical truth.

(This is known as the problem of *logical omniscience*.)

- Assuming that A only believes the two propositions above, there is a host of propositions that are *compatible* with A's beliefs. For example...
  - Since A is agnostic whether the NY Giants won the Super Bowl, there will be at least one $w' \in X_A$ such that the NY Giants won the Super Bowl in w', but also at least one $w'' \in X_A$ such that the NY Giants did not win the Super Bowl in w''.
- In general terms, any proposition logically consistent with either of the propositions that A believes is *compatible* with A's beliefs. I.e.

A proposition p is *compatible* with a set of worlds X if and only if $p \cap X \neq \varnothing$

13.1.2 Accessibility Relations

- An alternative way of capturing the relation between the agent's antecedent beliefs and the agent's reported belief is using so-called *accessibility relations*.
- Accessibility relations were originally introduced and used in modal logic where formal constraints on the accessibility between worlds yield different logical systems.
- An accessibility relation is a relation between possible worlds. For example, the expression '$R(w,w')$' says that w has access (or *can see*) the world w'.
- A semantics for beliefs can be formulated in terms of accessibility. For example, we might assume the following.

$$\ulcorner R^{\mathcal{B}_x}(w,w') \urcorner \text{ holds iff } w' \text{ is } \textit{compatible} \text{ with } x\text{'s beliefs in } w.$$

- We can then give the following lexical entry for 'believe'.

(313) $[\![\textit{believe}]\!]^{g,w} = \left[\lambda p_{\langle s,t\rangle}.\ [\lambda x_e.\ \forall w' \in \mathcal{W}\colon R^{\mathcal{B}_x}(w,w') \rightarrow p(w') = 1]\right]$

13.1.3 Formal Properties of Accessibility Relations

- Different attitude verbs will have different accessibility relations, but it is useful to consider some of the formal properties of these relations.

  · **Reflexivity**: $\forall x\big(R(x,x)\big)$

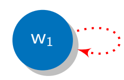

In set theoretic terms, a relation R is reflexive if and only if for every member x of the domain of R, R contains the ordered pair $\langle x,x\rangle$.

  · **Transitivity**: $\forall x\forall y\forall z\Big(\big(R(x,y) \wedge R(y,z)\big) \rightarrow R(x,z)\Big)$

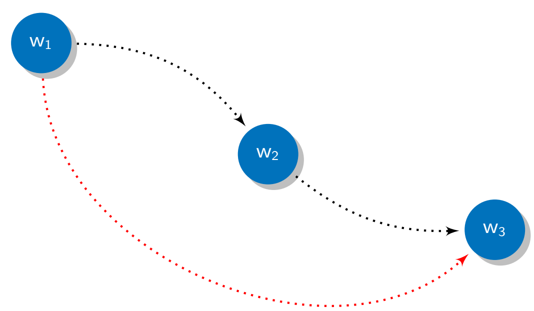

In set theoretic terms, a relation R is transitive if and only if for every ordered pairs $\langle x,y\rangle$ and $\langle y,z\rangle$, the pair $\langle x,z\rangle$ is also in R.

  · **Symmetry**: $\forall x\forall y\big(R(x,y) \rightarrow R(y,x)\big)$

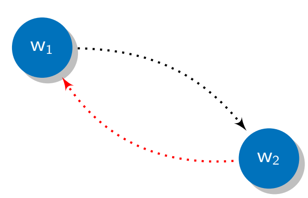

In set theoretic terms, a relation R is symmetric if and only if for every pair $\langle x,y\rangle$ in R, the pair $\langle y,x\rangle$ is also in R.

- Now consider the sentence below.

(314) Peter knows that grass is green.

- The propositional attitude verb 'know' is *factive*, i.e. for (314) to be true, the complement clause ('that grass is green') must be true.
- So, the accessibility relation associated with know must be *reflexive* — i.e. the set of worlds at which the complement clause is true *must* include the actual world.

(315) $[\![know]\!]^{g,w} = \lambda p_{\langle s,t\rangle}.\ [\lambda x_e.\ \forall w' \in \mathcal{W}\colon R^{\mathcal{K}_x}(w,w') \rightarrow p(w') = 1]$

- In contrast, the accessibility relation for 'believe' is non-reflexive (as opposed to irreflexive which is the opposite of reflexive, i.e. $R(w,w)$ is always false). For (309) to be true, it's not required that the complement clause is true in the actual world.
- In modal logic, '$\Box$' is used as a universal quantifier over possible worlds. So, '$\Box\phi$' essentially means that for every possible world w, ϕ is true at w.
- So, in modal logic, *reflexivity* corresponds to the following axiom:

(reflexivity)

$$\Box\phi \rightarrow \phi$$

- And, *transitivity* corresponds to the following axiom:

(transitivity)

$$\Box\phi \rightarrow \Box\Box\phi$$

- If the accessibility relation for 'know' is transitive, it would mean that if s knows that ϕ, then it follows that s knows that s knows that ϕ.
- Whether 'knowledge' satisfies this axiom is a controversial question. But there are other propositional attitude verbs for which transitivity seems to clearly fail. For example,

  s fears that $\phi \nrightarrow s$ fears that s fears that ϕ.

  s hopes that $\phi \nrightarrow s$ hopes that s hopes that ϕ.

- In modal logic, symmetry corresponds to the following axiom:

(symmetry)

$$\phi \rightarrow \Box\Diamond\phi$$

- If this axiom holds for knowledge, it implies that if something is the case at the actual world, then you know that *it is compatible* with what you know.
- **Equivalence Relations**
  In research on AI and Computer Science, the models used there are often assumed to be *equivalence relations*. An equivalence relation is a relation which is *reflexive*, *transitive*, and *symmetric*.
- Simplifying slightly, an equivalence relation guarantees that every world has access to every other world. A consequence of this is that the following axiom is valid:

(negative introspection)

$$\neg\Box\phi \rightarrow \Box\neg\Box\phi$$

- This axiom is called 'negative introspection' and it is a fairly implausible axiom for most propositional attitudes. For example, for knowledge, this axiom implies that if an agens s does not know that p, then s knows that s doesn't know that p. I.e. s can introspect what things s fails to know. It would be great if this was always true, but sadly it is not.

13.2 De Dicto and De Re

- In the lectures on quantifiers, it was noted that sentences with more than one QNP give rise to structural ambiguities. For example, (316) below is ambiguous between the two readings in (316a) and (316b).

(316) Every students wants a car.
  a. $\forall x\big(\text{student}(x) \rightarrow \exists y(\text{car}(y) \wedge \text{wants}(x,y)\big)$
  b. $\exists y\big(\text{car}(y) \wedge \forall x(\text{student}(x) \rightarrow \text{want}(x,y)\big)$

- However, a similar kind of ambiguity seems to arise when QNPs occur within a propositional attitude ascription. Consider the sentences below.

(317) Ralph believes that the president is a traitor.
(318) Ralph wants to marry a plumber.

- The sentence in (317) can be true relative to (at least) the following two potential circumstances:

  i. Suppose that Ralph believes, for some reason, that anyone who takes up the office of the presidency is per default a traitor. (317) then seems true regardless of whether Ralph knows who the current president is.
  ii. Suppose that Ralph believes that Barack Obama is a traitor, but that he is unaware that Obama is the president. In that case, (317) also seems true.

- Similarly, the sentence in (318) can be true relative to (at least) the following two sets of circumstances:

  i. Suppose that Ralph desires to marry any individual who is a plumber, but that he has no plans to marry anyone specific. In that case, (318) also seems true.
  ii. Suppose that Ralph wants to marry Louise and that, unbeknownst to Ralph, Louise is a plumber. In that case, (318) also seems true.

- The two types of interpretations displayed above are standardly referred to as 'de dicto' and 'de re' interpretations. For example, when the definite description in (317) is interpreted *de dicto*, the description itself is part of the content of the attitude in question. In contrast, when it is interpreted *de re*, it is not part of the content of the attitude.
- De dicto and de re ambiguities have standardly been thought to be structural ambiguities similar to the quantifier ambiguities.
- To illustrate, if we assume that 'the F' is a genuine QNP (thus type $\langle et,t\rangle$ and essentially an existential quantifier), this ambiguity can be represented in terms of the scope of the quantifier and the attitude verb:

(317) Ralph believes that the president is a traitor.
  a. $\text{BEL}_{\text{Ralph}}\big(\exists x(\text{president}(x) \wedge \forall y(\text{president}(y) \rightarrow y = x) \wedge \text{traitor}(x)\big)$

b. $\exists x\big(\text{president}(x) \wedge \forall y(\text{president}(y) \rightarrow y = x) \wedge \text{BEL}_{\text{Ralph}}(\text{traitor}(x))\big)$

- This would correspond to the difference between the following LFs.

(319)

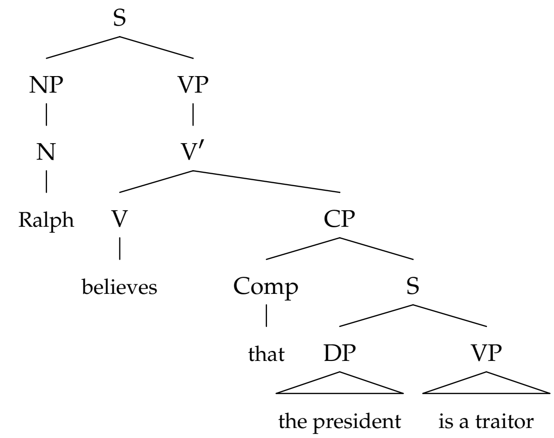

(320)

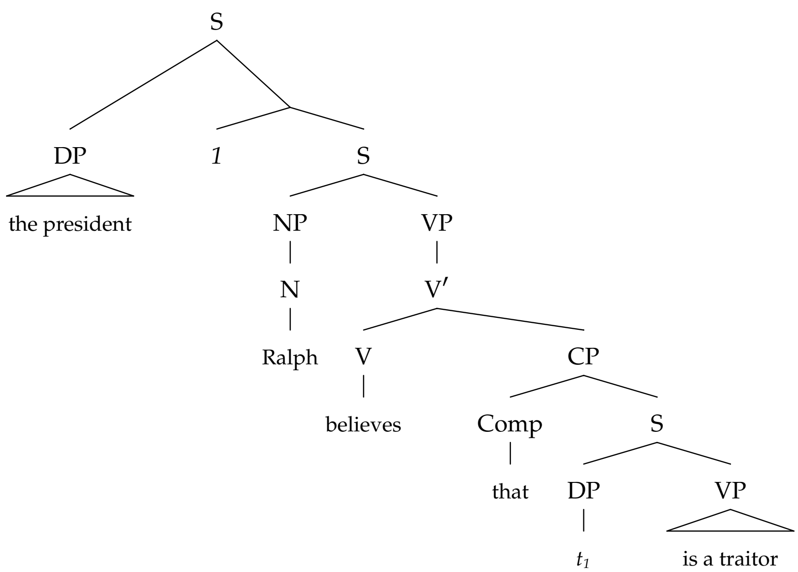

- Notice that the difference in scope between 'the *F*' and an attitude verb will be truth conditionally significant only if 'the *F*' is treated as an existentially quantified claim.
- By contrast, if 'the *F*' is a referential term (as proponents of the Frege-Strawson analysis maintain), the sentence has only one interpretation regardless of the scope of the description.
- Since definite descriptions do seem to give rise to these allegedly structural ambiguities, this is often considered a point in favor of the Russellian analysis of definite descriptions.

Chapter 14

Modality

14.1 Modals as Quantifiers over Possible Worlds

- In philosophy, the notions of *necessity* and *possibility* are standardly analyzed in terms of possible worlds. For example, it is standardly assumed that a proposition p is *possible* if and only if there is a possible world w where p is true. A proposition p is *necessary* if and only if p is true at every possible world w.
- We use '$\Box$' for *necessity modals* ('must', 'necessarily', 'have-to') and '$\Diamond$' for *possibility modals* ('might', 'may', 'could'). Necessity modals correspond to *universal* quantification over a set of possible worlds and possibility modals correspond to *existential* quantification.
- As a first pass, one might propose lexical entries such as these:

(321) $[\![must]\!]^{g,w} = [\lambda p_{\langle s,t\rangle}.\ \forall w'\!: p(w') = 1]$
(322) $[\![may]\!]^{g,w} = [\lambda p_{\langle s,t\rangle}.\ \exists w'\!: p(w') = 1]$

- Whether a modal makes a universal claim or an existential claim depends on its *modal force*. Above, we are assuming that there are only two options, namely universal force or existential force.

This might be inaccurate as there are modals such as 'should' and 'ought to' that are often labeled as *weak necessity modals*. These modals are generally considered stronger than possibility modals, but not quite as strong as necessity modals.

- Notice that the universal and existential quantifiers are duals, i.e.

(323) a. $\forall x\phi \leftrightarrow \neg\exists x\neg\phi$
b. $\forall x\neg\phi \leftrightarrow \neg\exists x\phi$

(324) a. $\exists x\phi \leftrightarrow \neg\forall x\neg\phi$
b. $\exists x\neg\phi \leftrightarrow \neg\forall x\phi$

- Hence, from this somewhat crude analysis, several desirable inferences come "for free". For example:

(325) a. You must leave → It's not the case that you may stay.
b. $\Box\phi \rightarrow \neg\Diamond\neg\phi$

(326) a. You may leave → It's not the case that you must stay.

b. $\Diamond\phi \rightarrow \neg\Box\neg\phi$

- We also get a straightforward explanation of the inconsistency of (327) and the consistency of (329).

(327) You must leave and you must stay. (contradiction)
(328) $\Box\phi \wedge \Box\neg\phi$

(329) You may leave and you may stay. (consistent)
(330) $\Diamond\phi \wedge \Diamond\neg\phi$

- With the lexical entries for 'might' and 'must' above, only *logical* necessities and possibilities seem to be modelled — i.e. things that are true at every world in W, i.e. literally *every* possible world.
- But, ideally, we would want to be able to predict that sentences such as (331) and (332) are, at least, sometimes true.

(331) Jack must go to jail for his crimes.
(332) Sue must have missed the train.

- Yet, with our current lexical entries, sentences such as (331) and (332) will never be true, since there clearly are *possible* worlds where Jack does not go jail for his crimes or where Sue did not miss the train in question.

14.2 Different Kinds of Modals

- One first step towards making our analysis of modals more flexible is to relativize the analysis to the actual world.
- For example, one could think that whether a sentence such as (331) is judged true depends on what current information is available. For example, suppose that Sue was supposed to be on the 5.00 o'clock train, but that she left her house late, and that as the train comes into the station, we discover that she is not on it. In that case, (332) seems true.
- In contrast, if we witnessed Sue board the train, then even if she doesn't seem to be on the train now, (332) seems false.
- So, we could try to change the lexical entries for 'may' and 'must' so that the only relevant possible worlds are those compatible with the available *information* (where 'information' here means something like evidence), i.e.

(333) $[\![may]\!]^{g,w} = \lambda p_{\langle s,t\rangle}.\ \exists w'$ compatible with the evidence in w: $p(w') = 1$
(334) $[\![must]\!]^{g,w} = \lambda p_{\langle s,t\rangle}.\ \forall w'$ compatible with the evidence in w: $p(w') = 1$

- While these lexical entries seem like a good start, it is not clear that they will work for other modalized sentences. Consider for example the following sentence.

(335) You must wear your seat belt.

- This sentence could both be true or false — it seems true when there is a rule or law which states that wearing a seat belt is required. In contrast, it seems false when there is no such rule or law.
- But since the truth of (335) depends on certain rules or laws — rather than any kind of evidence, we seem to need a lexical entry like the following:

(336) $[\![must]\!]^{g,w} = \lambda p_{\langle s,t \rangle}.\ \forall w'$ compatible with the rules in w: $p(w') = 1$

Regarding the notion of *compatibility*, we assume here that a world w' is *compatible* with the rules of w if and only if *nothing happens* in w' which violates the rules of w.

- In other words, we now seem to need at least two different lexical entries for 'must'. Now consider this third case.

(337) The button must be pushed after connecting the cable.

- Again, it is easy to imagine a context in which this sentence is true and a context in which it is false.
  - Suppose that the machine in question works only if the button is pushed after the cable is already connected.
  - Suppose that the machine in question works even if the button is pushed before the cable is connected.
- Notice that the truth of (337) does not seem to turn on either evidence, rules, or laws. Instead, the truth of (337) seems to depend simply on the constitution of the machine in question. So, this looks like yet another way of using the modal 'must'.

14.3 Modal Flavors

- What the examples above demonstrate is that modals have different *flavors* — i.e. modals can be used to make a variety of seemingly different claims.

— **Epistemic Modals**: Modals relating to bodies of knowledge or belief.

(338) It *might* rain today. (given what is believed/known)
(339) It *must* be raining. (given what I believe/know)

— **Deontic Modals**: Modals relating to sets of rules, laws, permissions, or obligations.

(340) Bob *could* face time in jail. (given the laws of the country)
(341) Bob *must* resign from his post. (given the rules of the university)
(342) Bob *should* apologize to his parents. (given the prevailing moral norms)

— **Teleological Modals**: Modals relating to sets individual goals and plans.

(343) Bob *must* kick the ball harder. (given his goal to score a goal)
(344) Bob *could* do the extra credit exercises. (given his goal to get the highest GPA)

— **Bouletic Modals**: Modals relating to e.g. desires and wishes.

(345) Bob *should* publish a book. (given his desire to be promoted)
(346) Bob *could* go to the dinner. (given his desire to be seen)

— **Ability Modals**: Modals relating to e.g. abilities.

(347) Bob *can* lift an elephant. (given his physical abilities)

- As should be clear, the same modal expression can, on different occasions, express different kinds of modality.
- According to Portner (2009), English modals have the following possible flavors.

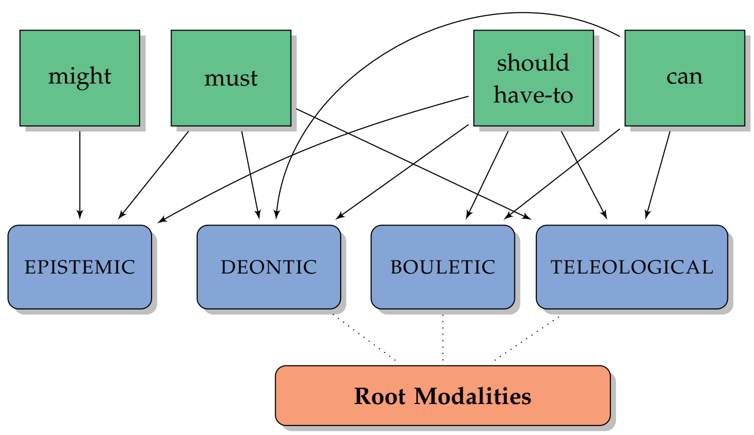

- This multiplicity of modal meanings is widely attested cross-linguistically, see e.g. Hacquard (2010) for discussion.
- This makes it very implausible that this is a genuine case of lexical ambiguity.

14.4 Kratzer's Analysis of Modals

- The predominant analysis of modal expressions today is the analysis put forward in a series of papers by the German linguist Angelika Kratzer. Kratzer's view is that the only meaning that modal expressions bring with them intrinsically is their force, i.e. whether they are universal or existential quantifiers over worlds.
- The domain over which the modals quantify is however context-sensitive. In other words, Kratzer's idea is that a modal sentence contains a hidden indexical element that is supplied by context. However, for modals this hidden indexical element is a variable ranging over sets of possible worlds (so, propositions).
- So, a sentence such as (348) which we have been assuming has the LF in (348a) actually has the LF in (348b)

(348) Bill must be tired.

a. [must [$_{\text{CP}}$ Bill is tired]]

b. [[must [$p_{\langle s,t\rangle}$]] [$_{\text{CP}}$ Bill is tired]]

14.4.1 Modals as proper quantificational determiners

- On our prior analysis ‘must’ was treated as a kind of generalized quantifier — it had semantic type $\langle st,t\rangle$
- On this alternative analysis, ‘must’ is now more like a quantificational determiner — it has the semantic type $\langle st,\langle st,t\rangle\rangle$
- The free variable in (348b) ranges over propositions, so its denotation is a (characteristic function over a) set of worlds.
- The idea then is that context can determine different sets of worlds over which the modal quantifies. For example:
  - The set of worlds compatible with what is known. (epistemic modality)
  - The set of worlds where the laws/rules are obeyed. (deontic modality)
  - The set of worlds where certain plans are realized. (bouletic modality)
- Modal sentences then have the following structure.

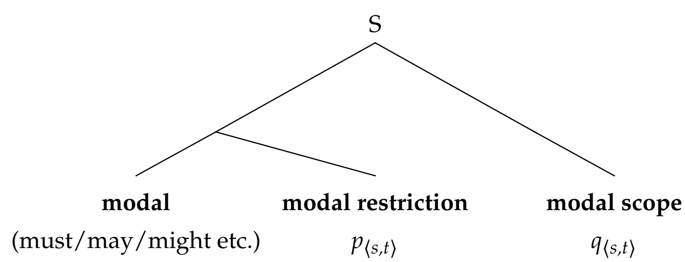

- This analysis enables us to provide just one lexical entry for modals with universal force and just one lexical entry for modals with existential force.

(349) $[\![\textit{must}]\!]^{g,w} = [\![\textit{have-to}]\!]^{g,w} = [\![\textit{necessarily}]\!]^{g,w} = \ldots =$
$[\lambda p_{\langle s,t\rangle}.\ [\lambda q_{\langle s,t\rangle}.\ \forall w'$ such that if $p(w') = 1$, then $q(w') = 1]]$

(350) $[\![\textit{can}]\!]^{g,w} = [\![\textit{may}]\!]^{g,w} = [\![\textit{might}]\!]^{g,w} = \ldots =$
$[\lambda p_{\langle s,t\rangle}.\ [\lambda q_{\langle s,t\rangle}.\ \exists w'$ such that $p(w') = 1$ and $q(w') = 1]]$

- A modal now expresses a relation between sets of worlds (as opposed to a relation between sets of individuals). This relation is either the subset relation (universal quantification) or non-disjointness (existential quantification).
- However, there is still a problem with these lexical entries — once again, modal sentences are predicted to be non-contingent.

- The solution is to make the restrictor sensitive to the world of evaluation.
- So, instead of thinking of the restrictor simply as a set of worlds (or rather as the characteristic function of a set of worlds, viz. a proposition determined by context), we should think of the restrictor as a function from a world (i.e. the evaluation world) to a proposition.
- For example, for the sentence (351), we assume the LF in (352).

(351) Bill must be quiet.
(352) [[must [$R_{\langle s\langle s,t\rangle\rangle}$]] [$_{\mathrm{CP}}$ Bill is quiet]]

- And then amend the lexical entries as follows:

(353) $[\![\mathit{must}]\!]^{g,w} = [\![\mathit{have\text{-}to}]\!]^{g,w} = [\![\mathit{necessarily}]\!]^{g,w} = \ldots =$
$[\lambda R_{\langle s,\langle s,t\rangle\rangle}.\ [\lambda q_{\langle s,t\rangle}.\ \forall w'$ such that if $R(w)(w') = 1$, then $q(w') = 1]]$

(354) $[\![\mathit{can}]\!]^{g,w} = [\![\mathit{may}]\!]^{g,w} = [\![\mathit{might}]\!]^{g,w} = \ldots =$
$[\lambda R_{\langle s,\langle s,t\rangle\rangle}.\ [\lambda q_{\langle s,t\rangle}.\ \exists w'$ such that $R(w)(w') = 1$ and $q(w') = 1]]$

- The relation R is a function from evaluation world w to a proposition (i.e. a set of worlds).
- This allows us to capture that when a modal has, e.g. an epistemic flavor, the set of worlds over which the modal quantifies is the set of worlds compatible with what is known *in the evaluation world*.

Chapter 15

Conditionals

15.1 Material Implication

- The simplest analysis of conditional statement — sentences of the form 'if ϕ, ψ' — is the truth functional analysis of classical logic known as *material implication*.
- On this analysis, a conditional statement is true if and only if either the antecedent is false or the consequent is true. I.e.

ϕ	ψ	$\phi \rightarrow \psi$
T	T	T
T	F	F
F	T	T
F	F	T

- Consider the sentence (355).

(355) If I'm healthy, I will come to class.

- An analysis of this sentence in terms of material implication seems to deliver quite reasonable results.
  - Consider the case where both antecedent and consequent is true: I'm healthy and I come to class. The conditional seems true (as predicted).
  - Consider the case where antecedent is true, but consequent is false: I'm healthy, but I don't come to class. I have broken my promise and the conditional seems false (as predicted).
  - Consider the cases where the antecedent is false: Since I have made no promises about what would happen in such a case, I may or may not come to class. In either case, I have not broken a promise, and so the conditional seems true (as predicted).
- We can implement this analysis by giving the following semantics for 'if'.

(356) $[\![if]\!] = [\lambda p_t \,.\, [\lambda q_t \,.\, p = 0 \text{ or } q = 1]]$

15.1.1 Problems with Material Implication

- Analyzing 'if' as a simple truth functional connective is riddled with problems.
- The most famous problems are the *paradoxes of material implication*. These "paradoxes" arise because this analysis assumes that the truth of a conditional statement *only* requires either that the antecedent be false or that the consequent be true.
- This means that the conditional statements below are all predicted to be true.

(357) If exactly 100 people live in China, exactly 2 people live in China.
(358) If Mitt Romney owns less than two cars, 2 + 2 = 5.
(359) If 2 + 2 = 5, then 2 + 2 = 4.
(360) If Goldbach's conjecture is true, then I'm a lecturer in philosophy.

- These are obviously quite bad predictions, but this is just the tip of the iceberg.
- When a conditional statement is analyzed in terms of a disjunction, the following equivalences are valid.

$$\begin{array}{lllll} \neg(\phi \rightarrow \psi) & \equiv & \neg(\neg\phi \vee \psi) & & \\ & & \neg(\neg\phi \vee \psi) & \equiv & \neg\neg\phi \wedge \neg\psi \\ & & & & \neg\neg\phi \wedge \neg\psi \quad \equiv \quad \phi \wedge \neg\psi \end{array}$$

- Given this, we now predict that (361) is true if and only if (362) is true.

(361) It's not true that if there will be a minor earthquake in Edinburgh tomorrow, the Dugald Stewart Building will collapse.

(362) There will be a minor earthquake in Edinburgh tomorrow and the Dugald Stewart Building will not collapse.

15.1.2 Material Implication and Embedded Modals

- The problems described above are exacerbated when the conditional statements contain embedded modal expressions.
- Imagine that we are lost, but after looking at a map, I assert (363).

(363) If we are on Route 183, we might be in Lockhart now.

- There are at least two possible LFs that we might consider for this sentence, namely one where the modal is raised to take wide scope over the whole conditional and one where the modal is only raised to take scope over the consequent.

(364)

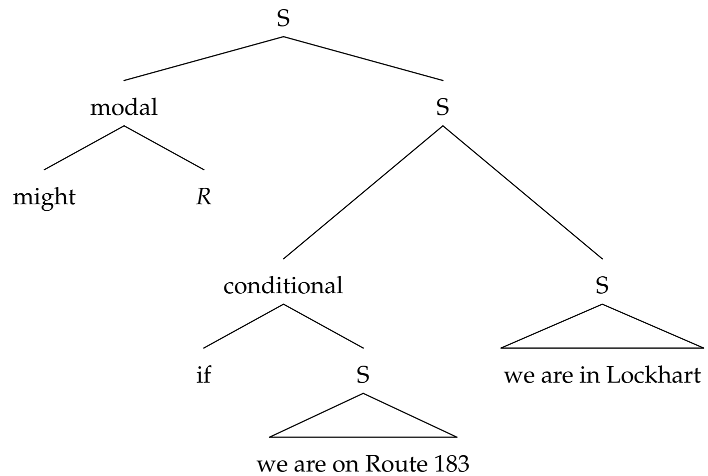

- (364) has the following structure.

$$might(\text{if } \phi, \psi)$$

- With this LF for (363), we predict that (363) is true if and only if either (a) or (b).

  (a) $\exists w'$ accessible from w where $[\![\textit{we are on Route 183}]\!]^{w'} = 0$

  (b) $\exists w'$ accessible from w where $[\![\textit{we are in Lockhart}]\!]^{w'} = 1$

- So, assuming that the accessibility relation (the modal base) is *epistemic*, then all it takes for (364) to be true is either (i) or (ii).

  (i) It is consistent with what we know that we are not on Route 183.
  (*if so, then there is possible world where the antecedent is false, and hence a possible world where the conditional is true — regardless of geographical facts*).

  (ii) It is consistent with what we know that we are in Lockhart
  (*if so, then there is a possible world where the consequent is true, and hence a possible world where the conditional is true — regardless of geographical facts*).

- Clearly, this makes the conditional as a whole come out true too easily.
- Let's consider the alternative LF.

(365)

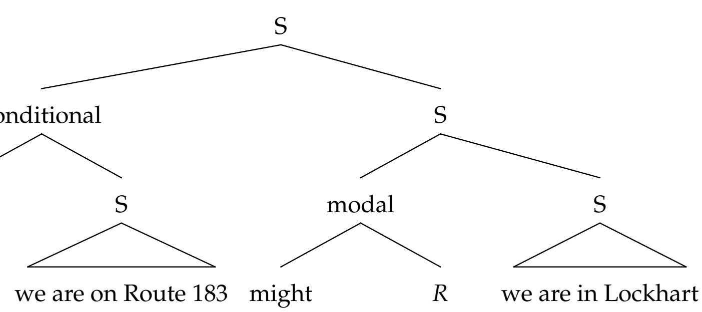

- (363) has the following structure.

$$\text{if } \phi, \textit{might } \psi$$

With this LF for (363), we predict that (363) is true if and only if either (a) or (b).

(a) $[\![\textit{we are on Route 183}]\!]^{w} = 0$

(b) $\exists w'$ accessible from w where $[\![\textit{we are in Lockhart}]\!]^{w'} = 1$

- If we assume that the accessibility relation (the modal base) is epistemic, then all it takes for (364) to be true is that either (i) or (ii) is true.

(i) We are not on Route 183.
(*if so, then the antecedent is false and as a result the conditional is true — regardless of geographical facts*).

(ii) It is consistent with what we know that are in Lockhart.
(*so, as long as our knowledge doesn't rule out that we are in Lockhart, the conditional as a whole is true — regardless of geographical facts.*)

15.2 Strict Implication

- One way to avoid (some of) these bad results is to give a modal analysis of 'if'.
- The strict implication analysis of conditional statements is basically a modalized version of material implication, namely material implication plus a necessity modal.
- A conditional statement of the form 'if ϕ, ψ' is then analyzed as follows:

$$\Box(\phi \rightarrow \psi)$$

- In short, on this modalized version, a conditional statement is true if and only if there is no possible world w where the antecedent is true and the consequent is false.
- If this analysis is chosen, our lexical entry for 'if' would look as follows.

(366) $[\![\textit{if}]\!]^{g,w} = [\lambda p_{\langle s,t\rangle}.\ [\lambda q_{\langle s,t\rangle}.\ \forall w'\text{: } p(w') = 1 \rightarrow q(w') = 1]]$

- Which is equivalent to (367) below.

(367) $[\![\textit{if}]\!]^{g,w} = [\lambda p_{\langle s,t\rangle}.\ [\lambda q_{\langle s,t\rangle}.\ \forall w'\text{: } p(w') = 0 \vee q(w') = 1]]$

- Of course, these lexical entries suffer from a familiar problem — there is no relativization to the actual world. We can fix this as we did earlier, by introducing a covert accessibility function (a covert modal restriction).

(368) $[\![\textit{if}]\!]^{g,w} = [\lambda R_{\langle s,\langle s,t\rangle\rangle}.\ [\lambda p_{\langle s,t\rangle}.\ [\lambda q_{\langle s,t\rangle}.\ \forall w'\text{: } (R(w)(w') = 1 \wedge p(w') = 1) \rightarrow q(w') = 1]]]$

- This analysis helps solve some of the most critical problems — for example, for (357), suppose we take the covert accessibility relation to be epistemic and suppose we have no idea how many people live in China — so it's consistent with what we know that e.g. 100 people live in China.

(357) If exactly 100 people live in China, exactly 2 people live in China

- This is now analyzed as follows:

  must(exactly 100 people live in China → exactly 2 people live in China)

- But since the antecedent logically entails the negation of the consequent, then as long as there is a possible world w' where the antecedent is true, the consequent is false at w'.
- Hence, (357) is predicted to be false — a big improvement over the simple material implication analysis.

15.2.1 Problems for Strict Implication

- The analysis of conditional statements as strict implications is however only a slight improvement over the simple material implication analysis.
- In particular, the strict implication analysis gives rise to problems that are quite similar to the paradoxes of material implication. For example, if the antecedent of some conditional statement is *necessarily* false (or simply inconsistent with the worlds accessible from the evaluation world), the conditional statement is predicted to be vacuously true!
- To see why, consider (357) again.
  - Since we *know* that more than 100 people live in China, there is no possible world w' consistent with our epistemic state where exactly 100 people live in China.
  - This means that at all the accessible worlds—viz. the worlds that satisfy $R(w)(w')$—$p(w') = 0$. So, one of the conjuncts in the antecedent in the formula below is false.

    $$\forall w' \colon (R(w)(w') = 1 \wedge p(w') = 1) \to q(w') = 1$$

    And this formula is equivalent to the formula below:

    $$\forall w' \colon \neg(R(w)(w') = 1 \wedge p(w') = 1) \vee q(w') = 1$$

  - Hence, when the antecedent is inconsistent with our knowledge (and the modal base is *epistemic*), the negated conjunction on the left hand side comes out true, and hence the whole disjunction comes out true! So, the conditional is predicted to be true.
  - But intuitively, it's not true—despite the fact that we know that more than 100 people live in China!
- A similar problem arises when the consequent of a conditional statement is necessarily true (or simply true at all the worlds accessible from the evaluation world).
- In those cases, the conditional is also predicted to be vacuously true regardless of the antecedent.

15.3 *If*-clauses as Restrictors

- The general idea (originating with David Lewis and implemented in detail by Angelika Kratzer) is to treat 'if' as a restrictor (rather than a truth functional connective) of the worlds over which a modal quantifies. In other words, and *if*-clause serves as a domain restriction over the worlds where the consequent must be true in order for the whole conditional to be true.

15.3.1 Covert Modals

- Kratzer's main assumption is that a sentence of the form 'if ϕ, ψ', is covertly modalized: 'if ϕ, ψ' is really 'must R(if ϕ, ψ)'.
- So, in order to evaluate the conditional for truth, first a modal base must be determined. The function of the *if*-clause is then simply to further restrict that modal base. Suppose that the modal base determined at the evaluation world is the set of worlds X. If so, the function of the if-clause is to further restrict that set to include only ϕ-worlds, viz. $X \cap \phi$.
- So, in a sentence such as (363) below.

(363) If we are on Route 183, we might be in Lockhart now.

- This conditional statement is covertly modalized, i.e.

(369) *must*(we are on route 183)(we might be in Lockhart).

- The modal requires a modal base and if we suppose that the modal base is epistemic, this determines the set of worlds compatible with our evidence at the evaluation world.
- The *if*-clause then further restrict this set to a set of worlds where we are on route 183 at every world and for the conditional to be true, it must then be the case that we *might* be in Lockhart at each of those worlds.

(370) $[\![\textit{must } R(\textit{if } \phi, \psi)]\!]^{g,w} = [\![\textit{must } R'(\psi)]\!]^{g,w}$ where $R' = (\{w' \mid R(w)(w')\} \cap \phi)$

- So,

(371) $[\![\textit{must } R(\textit{if } \phi, \psi)]\!]^{g,w} = [\![\textit{must}(\{w' \mid R(w)(w')\} \cap \phi)(\psi)]\!]^{g,w}$

- So, for a conditional of the form 'if ϕ, ψ' to be true, what must be true is that ψ is true at all the worlds in the intersection of the modal base (determined by the context) and ϕ.
- Hence,

(372) $[\![\text{if } \phi, \psi]\!]^{g,w} = 1$ iff $\forall v$ such that $v \in (\{w' \mid R(w)(w')\} \cap \phi)$, $\phi(v) = 1$.

References

von Fintel, Kai and Heim, Irene 2007. 'Intensional Semantics'. Unpublished ms.

Frege, Gottlob 1892. '*On* Sinn *and* Bedeutung'. In Beaney, Michael (ed.) *The Frege Reader (1997)*. Blackwell Publishing, pp. 151–172.

Grice, Paul 1989. 'Logic and Conversation'. In *Studies in The Way of Words*. Cambridge, Massachusetts: Harvard University Press, pp. 22–41.

Hacquard, Valentine 2010. 'Modality'. In Maienborn, Claudia, von Heusinger, Klaus and Portner, Paul (eds.) *Semantics: An International Handbook of Natural Language Meaning*. Berlin: Mouton de Gruyter (draft – forthcoming).

Heim, Irene and Kratzer, Angelika 1998. *Semantics in Generative Grammar*. Blackwell Publishing.

Jacobson, Pauline 2014. *Compositional Semantics: An Introduction to the Syntax/Semantics Interface*. Oxford Textbooks in Linguistics, Oxford: Oxford University Press.

Kamp, Hans 1981. 'A Theory of Truth and Semantic Representation'. In Groenendijk, Jeroen, Janssen, Theo M.V. and Stokhof, Martin (eds.) *Formal Methods in the Study of Language, Part 1*. Mathematical Centre Tracts, pp. 277–322.

Kaplan, David 1989. 'Demonstratives'. In Almog, Joseph, Perry, John and Wettstein, Howard (eds.) *Themes From Kaplan*, chap. 17. Oxford University Press, pp. 481–563.

Ladusaw, William 1979. *Polarity Sensitivity as Inherent Scope Relations*. Ph.D. thesis, University of Texas at Austin, Austin, Texas.

Lewis, David 1986. *On the Plurality of Worlds*. Blackwell.

Partee, Barbara H., ter Meulen, Alice and Wall, Robert E. 1990. *Mathematical Methods in Linguistics*. Dordrecht: Kluwer Academic Publishers.

Peters, Stanley and Westerståhl, Dag 2007. *Quantifiers in Language and Logic*. Oxford University Press.

Portner, Paul 2009. *Modality*. Oxford University Press.

Ross, John Robert 1967. *Constraints on Variables in Syntax*. Ph.D. thesis, MIT, Cambridge, Massachusetts.

Russell, Bertrand 1905. 'On Denoting'. *Mind*, 14: 479–493.

Strawson, Peter F. 1950. 'On Referring'. *Mind*, 59: 320–344.