

III Syntax

Syntax (from Greek *syn* (συν) ‘together’ and *taxis* (τάξις) ‘arrangement’ — i.e., ‘arranging stuff together’) is the branch of linguistic analysis that addresses the internal structure of SENTENCES and the PHRASES (aka CONSTITUENTS) out of which they are composed, or, put differently, the way in which phrases/constituents are arranged to form sentences.

III.1 The sentence: Subject, predicate, and tense

Like words (which we examined in segment II), sentences are universal ingredients of human natural language. And just as in the case of morphology, syntax faces a number of immediate questions about the object of study: What is a sentence? What is a phrase?

Some might define a sentence as ‘a string of words beginning with a capital letter and ending with a full stop’. But such a definition is obviously not a very helpful one: it defines the sentence in terms of arbitrary properties of (English) orthography.¹ There are plenty of languages that have no orthographies at all (signed languages, and languages of illiterate societies, where ‘illiterate’ means ‘lacking a written tradition’; the term often takes on a derogatory sense, which we will emphatically distance ourselves from); and even in languages that do have a regulated orthography, the orthographic rules by which these languages abide are typically imposed by an external authority (e.g., the Academy of Sciences). Neither the capitalisation rules of English nor its punctuation conventions are essential in any way to the natural human capacity to form sentences. So we are not going to define the sentence in such arbitrary, extraneous conventions as initial capitalisation and punctuation.

A more helpful approach to the sentence is to look at it from within, from the perspective of its essential component parts. All complete sentences of human natural language consist minimally of a SUBJECT and a PREDICATE.² In addition, complete sentences usually contain some marker of TENSE, expressing the time at which the proposition denoted by the sentence holds. In the very simple example in (1), which represents a sentence, the subject is *the students*, the predicate is *peruse the course material*, and the tense marker is *will*.³ The predicate denotes a property (‘studying the course material’), and that property is attributed to the referent of the subject, the group of people identified by *the students*. (We will talk more about the meaning of predication in segment IV.)

1 In all the linguistic examples presented in these lecture notes, we will refrain from using the standard orthographic conventions of an initial capital and final full stop, even when the examples represent complete sentences: linguistic examples are presented in isolation, not as contributions to a discourse (for which these conventions are intended).

2 It is important to restrict this statement to *complete* sentences. In natural conversations, we often find sentences that are incomplete. For instance, in reply to the question *What did you eat for lunch?*, someone could reply by saying simply *A sandwich* — a perfectly meaningful contribution to the conversation, pragmatically complete but syntactically apparently incomplete: on the surface, the reply only contains a subpart of the predicate; it lacks the verbal head of the predicate and the subject. Surface appearances to the contrary notwithstanding, ‘fragment answers’ such as *A sandwich* are arguably best represented in syntax as complete sentences, their incompleteness being the result of the fact that when they are phonologically interpreted, only a portion of the syntactic material is being pronounced. Such incomplete phonological spell-out of syntactic structures is referred to in the literature as ELLIPSIS.

3 In linguistic examples, even when they represent complete sentences, we will not use the standard orthographic conventions of an initial capital and final full stop.

(1) [Sentence [Subject the students] [Tense will] [Predicate peruse the course material]]

In the sentence in (2), we find plenty more words — but we are still dealing with just a single sentence, with just a single predicate and just a single subject. It just happens to be the case that in (2) both the subject and the predicate are quite a bit larger than they were in (1). But although the internal structure of the subject and the predicate is more complex in (2) than it is in (1), the global syntactic structure of the sentence in (2) is just the same as that of the sentence in (1).

(2) [Sentence [Subject the people enrolled in this seminar] [Tense will] [Predicate pay very close attention to the information in the lecture notes, the slide shows and the textbook]]

III.2 Phrases and categories

The collections of words enclosed inside the brackets labelled ‘Subject’ and ‘Predicate’ in (1) and (2) are called PHRASES — NOUN phrases in the case of *the students* and *the people enrolled in this seminar*; VERB phrases in the case of *peruse the course material* and *pay very close attention to the information in the lecture notes, the slide shows and the textbook*. The labels ‘noun’ and ‘verb’ name the part-of-speech categories to which the HEADS of the phrases in question belong: we call *the students* a noun phrase because *students* is a noun (a word that names an entity in the extra-linguistic universe); and *peruse the course material* is called a verb phrase because *peruse* is a verb (a word that names an event of some sort). The analyses of (1) and (2) given above use the FUNCTION of the phrases to label them; but it is more customary in syntax to label phrases by their CATEGORY, as in (1′) and (2′) (where we use the initials of ‘noun’ and ‘verb’ to label the categories), where each instance of a noun phrase (not just the subject *the students*) is identified by the label ‘NP’. The categorial labelling system allows us to identify two NPs in the first example; and in the second example, we find not just several NPs and two VPs, but also three PPs, two of which are headed by the preposition *in* and a third has *to* as its head, and an AP, whose head is the adjective *close*. Together, A, N, P and V constitute the set of LEXICAL CATEGORIES.

(1′) [Sentence [NP the students] [Tense will] [VP peruse [NP the course material]]]

(2′) [Sentence [NP the people [VP enrolled [PP in [NP this seminar]]] [Tense will] [VP pay [NP [AP very close] attention] [PP to [NP the information [PP in [NP the lecture notes], [NP the slide shows] and [NP the textbook]]]]]]]

An important reason why syntacticians prefer the labelling in (1′) and (2′) is that the category of a phrase always remains constant whereas its function can change in the course of the syntactic derivation. You can get an immediate taste of this by turning the sentence *the students will peruse the course material* into its passive form (more on this in §III.6), given in (3). This sentence still features *the students* as a phrase of the sentence, but it is no longer the subject of it: *the course material* (another noun phrase) has been syntactically promoted to the subject function, and *the students* now appears in a PP introduced by *by*.

(3) [Sentence [Subject the course material] [Tense will] [Predicate be perused by the students]]

(3′) [Sentence [NP the course material] [Tense will] [VP be [VP perused [PP by [NP the students]]]]]

III.3 Recursion

Just as there is no upper limit on the length of words (remember the discussion of *antidisestablishmentarianism* and *supercalifragilisticexpialidocious* in segment III), so sentences can be of infinite length. And just like complex words, sentences have the property of RECURSION: we can make them longer and longer by embedding a portion of structure of a particular kind inside a larger structure of the same kind. Indeed, sentences can themselves embed other sentences: in (4a), the relative clauses *who study in this department* and *which is called linguistics* are clauses (another name for ‘sentences’) that are an integral part of the sentence as a whole. The fact that sentences can be embedded in larger sentences can be illustrated also on the basis of examples of the type in (4b).

- (4) a. the people who study in this department study a subject which is called linguistics
 b. John thinks that Mary believes that Bill claimed that Bob said that Sue hoped that...

Sentences of the type in (4b) could in principle keep on growing indefinitely. Of course we would eventually lose track of the message that is supposed to be conveyed by such a complex sentence. But that is a restriction on the PROCESSING of sentences — not a restriction on the syntax of sentences. The syntax of sentences allows in principle for infinite RECURSION: the embedding of a structure of a certain type inside a larger structure of the same type.

We see recursion not just in the fact that sentences can be embedded in larger sentences. Phrases show recursion as well. Thus, the noun phrase *the students*, encountered previously, can embed another noun phrase, which can in turn embed another noun phrase, etc. etc.:

- (5) a. [_{NP} the students]
 b. [_{NP} the students [_{PP} of [_{NP} English]]]
 c. [_{NP} the students [_{PP} of [_{NP} the language [_{PP} of [_{NP} Shakespeare]]]]]
 d. [_{NP} the students [_{PP} of [_{NP} the language [_{PP} of [_{NP} the author [_{PP} of [_{NP} *Hamlet*]]]]]]]]]

From (5), we not only deduce that noun phrase embedding can be recursive: we also learn that PPs that are embedded inside noun phrases can themselves have other PPs embedded inside them. With prepositions, recursion can be illustrated in a particularly spectacular way: in (6b) and (6c), each preposition except for the last one (*under*) *directly* embeds another PP.


- (6) a. the dog crawled [_{PP} under [_{NP} the bed]]
 b. the dog crawled [_{PP} from [_{PP} under [_{NP} the bed]]]
 c. the dog crawled [_{PP} out [_{PP} from [_{PP} under [_{NP} the bed]]]]]

In (3'), we had already come across something similar to (6) in the realm of verbs: the auxiliary *be* (itself a verb) embeds the VP of the main verb, another case of immediate self-embedding recursion.

In (5) and (6), we arrived at recursive syntactic structures by embedding a phrase of a particular kind *below* a larger phrase of the same kind, making the right-hand node of the structure ever more complex. The tree structure for (6c) given in (6c') makes this graphically transparent.

- (6c')
-
- the dog crawled [_{PP} out [_{PP} from [_{PP} under [_{NP} the bed]]]]]

But when the syntax tries to produce a recursive structure in which ‘the middle’ of the tree keeps getting bigger, the parser quickly loses track when trying to make sense of the result. The sentence in (10) is already quite difficult to interpret; but when we leave out the instances of *that* which introduce each relative clause, as in (10’), the parser basically throws in the towel. The examples in (10) and (10’) illustrate so-called centre-embedding — an instance of recursion whereby the tree expands ‘in the middle’ rather than just on the left-hand or right-hand side. The syntax can create such structures: (10) and (10’) are well-formed. But the sentence processing mechanism has quite a bit of trouble with them, so it is unlikely that you will come across such constructions a lot in real life. This illustrates again that language is not just a matter of COMPETENCE — though we have the linguistic competence to handle multiple centre-embedding, in PERFORMANCE we struggle with it.



- (10) the rat that the cat that the dog chased bit died
 (10’) the rat — the cat — the dog chased bit died

III.4 Phrase structure and binary branching


At this point, let us return to (1’), repeated below. Here, the verb *peruse* and the NP of *the course material* together form the VP. The verb is the HEAD of the VP but it is not itself ‘the verb phrase’: it needs an object to complement it.

- (1’) [Sentence [NP the students] [Tense will] [VP peruse [NP the course material]]]

In (6a), we see that an NP can also be used as a COMPLEMENT to a preposition — and again, it is essentially inevitable to combine the preposition with a nominal complement: by itself, *under* would not survive here (though in *the ship went under*, it does).



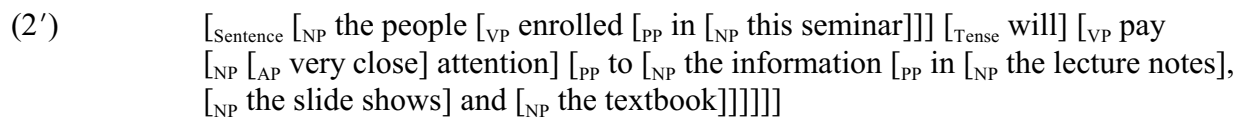
The way in which the verb *peruse* combines with *the course material* in (1’) is precisely the same as the way in which the preposition *under* combines with *the bed* in (6b). In both cases, the head combines with its complement in a SISTERHOOD relation, the two together forming a MOTHER NODE whose label is that of the head. Sisterhood is the relationship *par excellence* for combining material into ever larger syntactic PHRASE STRUCTURES.

In syntactic structures, it is usually the case that a mother node has at most two daughters — or, put differently, that sisterhood is confined to two elements. But for the structure of the sentence in (1’), this does not appear to be the case: in (1’), the NP *the students*, the Tense marker *will*, and the VP *peruse the course material* are three daughters of the same mother node, identified by the label ‘Sentence’ (or ‘S’, for short). The structure in (1’) was the standard representation of the syntax of the sentence for several decades. But it is strictly speaking out of date. More recent work has developed the structure of the sentence in such a way that it is no longer an exception to the generalisation that mother nodes have at most two daughters: the Tense marker first takes the VP as its complement, and the node thus formed combines with the subject to form the sentence. In this light, we can update (1’) as in (1’'), with an additional left bracket (which, to keep things simple, we will leave unlabelled here) between the subject and Tense, as a reflex of the hypothesis that Tense combines with VP before the subject comes into the picture. The bracketed structure in (1’') translates into the tree given above it.



The structure in (1'') faithfully reflects the BINARY BRANCHING hypothesis — the idea that syntactic structures (and presumably linguistic structures across the board, also in morphology, phonology, and semantics) strictly involve mother nodes that have at most two daughters. The binary branching hypothesis is not as such a fact: it is an assumption made about the structures tolerated by the human language faculty. We should certainly like it to be true — if it is, this means that children acquiring their native language have far fewer types of structures to contemplate for the utterances that they are presented with than they would if a mother node could in principle have any number of daughters. The binary branching hypothesis is a RESTRICTIVE hypothesis: it narrows down the analytical playing field, thereby easing the task faced by the child (and, into the bargain, the task faced by the linguist as well).

With the structure of the sentence in (1'') saving us from a collision with binarity, one more worry that needs to be addressed in this connection is the syntax of coordination. Consider again the structure in (2') (repeated below), in which the complement of the second token of the preposition *in* is a string of three NPs, the first two joined together by a comma and the second and third NPs separated by the conjunction *and*. If we analysed coordination in such a way that the three NPs and the conjunction are all on equal footing, we would arrive at a mother node with four daughters, in flagrant disregard of the binary branching hypothesis. This is shown in (11). There is an alternative to (11), however, which allows us to steer clear of trouble: in (11'), we treat both the conjunction and the comma in a way similar to the treatment of the Tense marker in (1''), as heads that take what follows them as their complement and then combine the node formed via complementation with a constituent to its left to form the complete coordination phrase ('CoordP', aka '&P').



Note that once coordination is represented as in (11'), we can chalk up multiple coordination constructions of the type illustrated in (2') as yet another illustration of the 'Russian doll' effect that is the essence of recursion: again, we see a tree in which a structure of a particular type (here, CoordP1) is embedded within a larger structure of the same type (CoordP2). The analysis in (11') also makes us see that coordination and subordination are not fundamentally different ways of putting things together in syntax — indeed, the syntactic processes by which coordination and subordination structures are formed are very much the same.

III.5 Phrase-structural ambiguity

Now that we have a basic sense of the workings of phrase structure, let us take a close look at a few concrete examples and determine their syntax. We will discover in this section that this task is quite often more formidable than we might at first have thought.

Consider first the sentence in (12), a classic example from Noam Chomsky's (1965) book *Aspects of the theory of syntax*. Though you might not realise this right away, this sentence is two-ways ambiguous, as indicated by the paraphrases given in (12a) and (12b). On the former reading, *flying* serves as a modifier of *planes*, and *flying planes* is a nominal constituent, an NP. On the interpretation in (12b), by contrast, *flying planes* is a kind of clause (traditionally called a 'gerund'), replaceable with an infinitival clause.

- (12) flying planes can be dangerous
a. 'planes that fly can be dangerous' [_{NP} flying planes]
b. 'to fly planes can be dangerous' [_S flying planes]

As soon as we remove the modal auxiliary *can* from these sentences and force the copula to inflect for the number specification of the subject, the ambiguity seen in (12) falls away: (13a) (in which *are* agrees in number with *planes*) only supports a reading in which *flying planes* is an NP; (13b) does not support this reading, and instead forces *flying planes* to be interpreted as 'to fly planes'.

- (13) a. flying planes *are* dangerous
b. flying planes *is* dangerous

The sentences in (13a,b) are both grammatical but each of them is unambiguous. The fact that the modal auxiliary *can* shows no agreement in number with its subject is responsible for the fact that (12) supports two different interpretations. On both of these readings, *flying planes* is a constituent serving as the subject of the sentence as a whole. The difference lies in the internal structure of this constituent: in the a-reading, *planes* is the head of *flying planes*, and *flying* is a modifier; in the b-reading, *flying* heads the phrase, and *planes* is its direct object. The difference between the two structures given to the right of (12) may look minimal; but the labelling of the nodes makes a very big difference in the external syntactic behaviour of the constituents in question.

Next, let us examine the sentence in (14), another classic from Chomsky's work:

- (14) John watched the man with binoculars
a. 'John watched the man who had binoculars'
 [_{VP} watch [_{NP} [_{NP} the man] [_{PP} with binoculars]]]
b. 'John used binoculars to watch the man'
 [_{VP} watch [_{NP} the man] [_{PP} with binoculars]]

Again, you might not realise right away that this sentence is ambiguous. The paraphrases in (14a) and (14b) help bring out the ambiguity. As in the case of (12), the ambiguity of this sentence is not an accident: it arises thanks to the fact that the sentence supports two different phrase-structural analyses — one in which *the man with binoculars* is a phrase or constituent (as shown in the structure below the paraphrase in (14a)), and one in which it is not (given below (14b)).

Just as in the case of (12), the difference between the two readings of (14) may seem quite minimal: in the bracketings given in (14) it boils down just to the absence or presence of a second ‘ $[_{NP}]$ ’ bracket to the right of *watch*. But this makes a world of difference: in the structure below (14a), the presence of this second ‘ $[_{NP}]$ ’ bracket indicates that the PP *with binoculars* is a subconstituent of the noun phrase of which *man* is the head; the absence of the second ‘ $[_{NP}]$ ’ bracket in the structure below (14b), by contrast, indicates that the PP *with binoculars* is not a constituent of the noun phrase of *man* but instead an immediate constituent of the verb phrase (VP).

This way of thinking about the ambiguity of (14) makes a straightforward prediction. Since in (14a) *with binoculars* is part of the complex noun phrase in object position, we expect that on this reading it should be possible to replace or syntactically displace the entire substring *the man with binoculars* as a constituent, by substituting a pronoun (*him*) for it or by putting it to the left of the subject (an operation called ‘topicalisation’) or promoting it to the subject function (via ‘passivisation’, which we already encountered briefly in (3’)). On the reading of (14) rendered in (14b) and analysed as in the structure underneath (14b), on the other hand, we expect such re- or displacement to fail; instead, we should only be able to re- or displace *the man*, leaving *with binoculars* behind. These expectations are fulfilled: (15a,b) are not ambiguous in the way that (14) is, supporting only the reading paraphrased in (14a); for (16a,b), by contrast, the reading in (14b) is readily available and it is very difficult to get the interpretation in which the binoculars are carried by the man.

- (15) a. the man with binoculars, John watched
 b. the man with binoculars was watched by John
- (16) a. the man, John watched with binoculars
 b. the man was watched with binoculars by John

Although the linear string in (14) does not ‘wear its syntax on its sleeve’, we can reliably bring its phrase structure to light with the help of syntactic constituency tests such as topicalisation and passivisation. The same is true for other cases of phrase-structural ambiguity. Consider (17):

- (17) Mary saw the boy walking towards the station
 a. ‘Mary saw the boy who was walking towards the station’
 b. ‘Mary saw the boy as he was walking towards the station’
 c. ‘Mary saw the boy as she was walking towards the station’
 d. ‘Mary witnessed the event of the boy walking towards the station’

This sentence is four-ways ambiguous, as the paraphrases in (17a–d) show. But readings (17b), (17c) and (17d)⁴ can all be eliminated by placing the substring *the boy walking towards the station* to the left of the subject, via topicalisation, or in the structural subject position, via passivisation: (18a,b) only have reading (17a). If, alternatively, we take just *the boy* and topicalise it or make it the subject of a passive sentence, as in (19), the interpretation paraphrased in (17a) is not available at all — precisely because *the boy walking towards the station* now is not a constituent.

- (18) a. the boy walking towards the station, Mary saw (but the boy in the car, she didn’t)
 b. the boy walking towards the station was seen by Mary

4 On this reading, *walking* can be replaced with the bare infinitive — *Mary saw the boy walk towards the station*.

- (19) a. the boy, Mary saw walking towards the station
 b.i the boy was seen walking towards the station by Mary
 b.ii the boy was seen by Mary walking towards the station

The two passive sentences in (19) differ with respect to where the *by*-phrase is placed, which appears to influence the reading most saliently available for these sentences: (19b.i) shows a preference for (17b) while (19b.ii) prefers (17c). For a proper understanding of this, we would need to probe into the nitty gritty of the syntax of passivisation — something for which we will not have an opportunity now. But such details aside, the discussion of (17) reaffirms the dependence of the interpretation of sentences on their internal structure. Semantic interpretation is not read off linear strings but off syntactic structures. The passive transformation helps us diagnose syntactic structure whenever the linear string does not wear its structural analysis on its sleeve.

III.6 Active and passive sentences: The passive transformation

The b-examples in (15), (16), (18) and (19) represent what we call PASSIVE sentences. The PASSIVE TRANSFORMATION is a rule of syntax that converts a construction with the signature to the left-hand side of the arrow in (20) into a construction in which the relative positions of the two NPs are swapped, and an auxiliary (*be*), participial morphology ('PTC') and a token of the preposition *by* are introduced.

$$(20) \quad \text{Passive Transformation} \\ \text{NP}_1 - \text{V} - \text{NP}_2 \quad \Rightarrow \quad \text{NP}_2 - \text{be} - \text{V} + \text{PTC} - \text{by NP}_1$$

Passivisation is a useful constituency test because it affects NP_2 as a constituent and promotes it to subject, and affects NP_1 as a constituent by demoting it to the complement of the preposition *by*. Yet, despite its generality and usefulness, the passive transformation is not applicable to all sentences with the signature ' $\text{NP}_1 - \text{V} - \text{NP}_2$ '. Take, for instance, the sentence *Mary became a successful scientist*, which has precisely this form. The passive sentence **a successful scientist was become by Mary* is ungrammatical. Unsuccessful passivisations such as this one have in the past been taken to argue against the existence in syntax of a general rule of passivisation. This argument is not valid, for two cogent reasons. First, there is really no alternative to a syntactic treatment of passivisation: syntax is inexorably implicated in the analysis of the passive because only syntax can manipulate the placement of phrases in sentences.⁵ And secondly, the ungrammaticality of **a successful scientist was become by Mary* is not in fact something that a syntactically oriented account of passivisation fails to make sense of: it is not some haphazard, mysterious exception that throws a wrench into the idea that passivisation is a general rule of syntax; it is a *principled* exception — an exception that proves the general rule, if you will.⁶

5 That passivisation cannot be treated as a lexical rule (more specifically, as an operation on the lexical argument structure of the verb undergoing the rule) is particularly clear from the grammaticality of (21b) in the text below, where the NP promoted to subject is not an argument of the passivised verb.

6 In general, exceptions by themselves never refute an analysis. Chomsky (1962/4) put it, 'Until incorporated in an explicit generative grammar, such examples simply stand as exceptions, no more relevant to the correctness of the already formulated rules than strong verbs and irregular plurals. Listing of innumerable examples is neither difficult nor very interesting; it is quite another matter to find rules that account for them, or a general theory of such rules.' What

What the structural description of the input to the passive transformation should take into account is whether NP₂ is an ARGUMENT or a PREDICATE. In *Mary invited a successful scientist*, NP₂ is an argument of the verb *invite*; in *Mary became a successful scientist*, by contrast, NP₂ does not refer to some entity in the universe outside of Mary but instead names a property that is attributed to Mary — NP₂ here is not an argumental expression; it is instead a predicate. We know this from the fact that in *Mary became a successful scientist*, it is possible to replace the verb with a token of the copula *be*, which always mediates a relationship between a subject and its predicate: indubitably, *Mary was a successful scientist* is a case in which the postcopular noun phrase is a predicate of *Mary*. Predicates are not eligible for promotion to subject under passivisation; only arguments are. This is something we need to add to the structural description of the passive rule: as it stands, (20) does not make this explicit. But once we make this small (but significant) amendment, the passive rule in (20) is safe as a rule of syntax.

It is important to note that although the NP promoted to subject by the passive rule must be an argument, it is *not* the case that the argument that undergoes promotion must be an argument of *the verb that is passivised*. We see this particularly clearly in sentences of the type in (21):

- (21) a. they proved the theory wrong
b. the theory was proved wrong

In (21a), they did not prove the theory — in fact, they did exactly the opposite: they proved that the theory was wrong. So in (21a), *the theory* is not an argument of *prove*; instead, it is the argument of *wrong*. Yet despite the fact that *the theory* is not an argument of the lexical verb that is passivised, this noun phrase happily undergoes promotion to subject in (21b).

The grammaticality of (21b) is one case that tells us in unmistakable terms that the passive transformation cannot be thought of as a lexical operation on the argument structure of the verb that is turned into a participle. Passivisation is not an operation on semantic structures — it is a syntactic transformation. It is not the case that passive sentences necessarily correspond to active sentences in which the subject is in control of the event — the ‘agent’, as semanticists say. In a sentence such as (22a), the subject is clearly not controlling the event: noises have no control over anything. Yet, (22a) can be passivised, as shown in (22b).

- (22) a. the noise frightened John
b. John was frightened by the noise

Not only is it not the case that the referent of the NP in the *by*-phrase of the passive has to be an agent (the example in (22b) told us that it does not have to be), it is actually possible for the subject of a passive sentence to be an agent. We see this in (23): although the referent of *John* is not in control of the *be led* event, he does have control over leaving the room — *John* is the agent of *leave the room*, yet appears in the subject position of a passive sentence.

- (23) a. the noise led John to leave the room
b. John was led to leave the room by the noise

is needed is an understanding of the exceptions. And typically, what one finds is that, once an understanding of the exceptions emerges, the exceptions in fact confirm the existence of a general rule.

The term ‘passive sentence’, which has firmly established itself in descriptive, educational, and theoretical linguistics alike, is unfortunate in making an unveiled appeal to a semantic intuition that holds for many pairs of active and passive sentences: the referent of the subject of the sentences in (15b) and (16b) (i.e., the NP of *man*) is not in active control of the event expressed by the verb phrase; instead, the subject plays a ‘passive’ role, undergoing the event, while the NP whose referent does have control over the event (*John*) appears not as the subject but as the complement of the preposition *by* (and can often remain entirely unexpressed: passive *by*-phrases are typically optional). An important message that the examples in (22) and (23) bring us is that there is no direct relationship between ‘semantic subject’ (or agent of the event) and ‘syntactic subject of the sentence’, or between ‘semantic object’ (or undergoer of the event) and ‘syntactic object of the sentence’. Which semantic argument is the syntactic subject is determined by properties of the syntax of the sentence, not by the semantics.

III.7 The structural subject

In light of the immediately preceding discussion, the STRUCTURAL SUBJECT of the sentence should not be defined in semantic or meaning-related terms. Being the structural subject of the sentence is, as the term suggests, a structural property, correlated with certain morphosyntactic characteristics.

The structural subject of the sentence is the argument which determines (in the simple present tense) whether we should use the third person singular form of the verb or not: only when the structural subject is third person singular does the verb take the morpheme *-s*, the marker of third person singular present-tense inflection (recall segment II). We see this in (24) and (25):

- (24) a. the students study/*studies the English language
- b. the student studies/*study foreign languages
- (25) a. the English language is/*are studied by the students
- b. foreign languages are/*is studied by the students

In (24a), where the subject (*the students*) is plural and the object (*the English language*) is singular, we get no *-s* form of the verb. The fact that *the English language* is singular does not give the verb the licence to be inflected with *-s*. In (24b), we do get an *-s* on the verb. The fact that the object of (24b) (*foreign languages*) is plural does not make this impossible: the object has nothing to say about the inflectional form of the verb; it is the fact that the subject of (24b) (*the student*) is singular that determines that the verb must be adorned with *-s*. In (25) we see that in the passive, it is the noun phrase of *language(s)* that is responsible for whether *-s* appears on the verb or not. Once again, it is the structural subject that is in control. The semantic role of the noun phrase is entirely irrelevant when it comes to the question of whether *-s* should appear on the verb or not: this decision is based on the morphological properties of the structural subject, regardless of its semantic role.

Another property of the structural subject, in languages such as English, is that it is the NP that inverts with the finite auxiliary in the formation of root questions. Imagine that instead of making the assertion in (26a), we want to ask whether it is true or not that the man with binoculars is eating sushi. The way to form the question, if we want to use it as a complete utterance by itself, is to place the finite auxiliary (*is*) to the immediate left of the subject (as a whole: not just to the left of a portion of it), as in (26b). This is called SUBJECT–AUXILIARY INVERSION (often abbreviated as ‘SAI’): the relative order of the subject and the finite auxiliary is turned around or inverted.

- (26) a. the man with binoculars is eating sushi
 b. is the man with binoculars eating sushi?
 *the man is with binoculars eating sushi?

Importantly, the syntactic operation responsible for the formation of root questions (*a*) inverts the relative order of the finite auxiliary and the structural subject regardless of the semantic role that the structural subject plays (i.e., the rule is insensitive to agenthood or other semantic attributes) and (*b*) can never change the relative positions of the finite auxiliary and a non-subject, not even if the non-subject is actually farther to the left than the structural subject. The examples below illustrate:

- (27) a. the noise is annoying the man
 b. is the noise annoying the man?
- (28) a. the man is being watched by John
 b. is the man being watched by John?
 *the man being watched by John is?
- (29) a. sushi, the man is eating
 b. *is sushi, the man eating?

III.8 Question formation and constituency

Both the passive transformation and question formation involve rules of syntax that serve as useful constituency tests. We had already seen this for the passive. Question formation can also help us out when we find ourselves in doubt as to the constituency of a sentence that we are presented with.

The diagnostic benefits of the syntax of question formation come to the fore in a number of different ways. Consider first the example in (30). The rule of SAI is defined in such a way that it takes the finite auxiliary and inverts it with the structural subject in its entirety. In (30), there are two finite auxiliaries, both spelled and pronounced the same way, as *can*. Which of these should we choose in the formation of the yes/no-question corresponding to (30)? The answer is perfectly clear to native speakers of English — in fact, it is clear to children who acquire English as their native language from the moment they are capable of forming questions with SAI. The fact of the matter is that (31a) is grammatical and (31b) is word salad.

- (30) people who can speak Portuguese can often speak Spanish as well
- (31) a. can people who can speak Portuguese __ often speak Spanish as well?
 b. *can people who __ speak Portuguese can often speak Spanish as well?

The reason why not even the youngest native-speaker producers of English root yes/no-questions make any mistakes with this is simple: (31b) fails to invert the finite auxiliary with the structural subject integrally; rather, it takes *can* out from within the structural subject and places it to its left, in flagrant violation of the SAI rule of English. The fact that (31b) is woeful tells us that [*people who can speak Portuguese*] is a constituent in (31) — more particularly, it is the constituent that occupies the structural subject position. It also tells us that when we say that SAI takes the finite auxiliary that is closest to the beginning of the sentence and places it in initial position, by ‘closest’ we mean ‘structurally closest’, not ‘linearly closest’.

Question formation is also diagnostically powerful for determining the constituency of strings that do not serve as subjects. But to see this, we need to move from yes/no-questions to questions in which we want to ascertain the nature of a particular element in the sentence, not the veracity of the sentence as a whole. Imagine that you can see that John is eating something, but you cannot quite figure out what it is that he is eating. In such a circumstance, you can ask someone who has a clearer view of John's plate:

(32) what is John eating?

The expression *what* (called a *wh*-expression because of the fact that (a) it is introduced by the letter combination *wh*- and (b) it introduces a question) is a constituent of the sentence that it occurs in — and tellingly, it occurs in a position in which one would not normally expect the direct object of *eating* to occur. English is a language in which the direct object usually follows the transitive verb. So we say *John is eating sushi*, not **John is sushi eating* or **sushi is John eating*.⁷ But when the direct object is being questioned, it shows up in sentence-initial position, immediately followed by the finite auxiliary in root questions (the latter as a result of SAI). The *wh*-phrase *what* has been displaced from its usual direct object position to sentence-initial position by the rule of *WH-FRONTING*.

The rule of *wh*-fronting takes the entire object, integrally, and places it in sentence-initial position. It cannot take just a subpart of the object and leave the rest behind. We see this in (33)–(35). From the ungrammaticality of the b-examples we learn that *a plate* in *a plate of sushi* (33), *sushi* in *Korean sushi (aka kimbap)* (34), and *your/whose* in *your/whose sushi* (35) are not constituents.⁸

- (33) a. John is eating a plate of sushi
 b. *what is John eating of sushi?
- (34) a. John is eating Korean sushi (aka kimbap)
 b. *what is John eating Korean (aka kimbap)?
- (35) a. John is eating your sushi
 b. *whose is John eating sushi?

The rule of *wh*-question formation as a constituency test is not just useful for objects: it applies equally well to the subject. Thus, consider (36) (a variation on the theme of (35)) and (37):

- (36) a. whose father is eating sushi?
 b. *whose is father eating sushi?

7 Topicalisation is an exception to the general rule that a non-*wh* direct object immediately follows the verb in English (put differently, that English is a 'VO' [V before object] language): (27a) is the product of a syntactic operation which in certain ways is similar to *wh*-fronting yet in another way is markedly different — it does not give rise to SAI.

8 The non-constituency of *your/whose* in (35b) this may come as a surprise, but the conclusion is arguably correct — notwithstanding the fact that *whose* sometimes does appear to serve as a constituent all by itself (as in *whose is this?*). Readers familiar with Hungarian will note that in this language, dative-marked possessors can legitimately be *wh*-fronted leaving the rest of the object noun phrase behind: *kinek ismered a fiát?* '(lit.) whose do you know the son, i.e., whose son do you know?'. The syntax of Hungarian and English diverges on this point — by no means the only point of variation between the two languages.

- (37) a. the man {in the black hat/who is wearing a black hat} is eating sushi
 b. *who is {in the black hat/who is wearing a black hat} eating sushi?

Although *the man* certainly can be a constituent when it occurs unaccompanied by modifier, we see that when *the man* is a subject modified by a PP or relative clause, it cannot by itself be placed to the immediate left of the finite auxiliary in a root *wh*-question. This is because in the strings *the man in the black hat* and *the man who is wearing a black hat*, the substring *the man* (i.e., the combination of the definite article and the head noun) is not a constituent. This tells us something very meaningful about the syntactic structure of complex noun phrases such as *the man in the black hat* and *the man who is wearing a black hat*. But we will save the investigation of this for a more advanced course in syntax.

III.9 The importance of abstraction from the surface

Having syntactic diagnostics (such as passivisation and question formation) that help us probe into the structure of sentences is useful because the linear string of words in a sentence frequently hides from view the dark secrets of the syntactic structure of the sentence.

To see this, we will examine first the sentence in (38). You will hit upon a particular reading for this sentence right after you finish reading it — and that reading will likely blind you temporarily for other possible interpretations. But there is in fact more than one reading available for (38). The paraphrase in (38a) treats *someone* as the subject of *help out*: whoever this person may be, (s)he takes it upon him/herself to help out some unspecified people. By contrast, in (38b) *someone* is construed as the object of *help out*, receiving rather than giving help.

- (38) there is always someone to help out
 a. ‘there is always someone that can help people out’
 b. ‘there is always someone that people need to help out’

The ambiguity of (38) disappears as soon as we provide an explicit object or subject for *help out*: in (39a), *someone* must be the helper; in (39b), it must be the ‘helpee’.

- (39) a. there is always someone to help *them* out
 b. there is always someone for *them* to help out

The non-ambiguity of (39a,b) is unsurprising: in each of these sentences, there is only one grammatical function that *someone* could have within the infinitival clause; the other function is already ‘taken’ by *them*. Much more noteworthy is the fact that the surface syntax of (38) underdetermines the grammatical function of *someone* in the infinitival clause, and that as a result, this single string is amenable to two interpretations, differing quite dramatically.

How do we structurally bring out the difference between the two readings of (38)? The NP of *someone* is associated with an infinitival relative clause, *to help out* — *someone* itself stands outside this relative clause, but it is associated with a ‘placeholder’ for it inside the relative clause. Let us call this placeholder ‘PRO’, a silent pronoun. This pronoun can be mapped either into the subject position or into the object position of the infinitival relative clause, as in (40a) and (40b), resp. (where we use coindexation to establish the link between *someone* and ‘PRO’).

- (40) a. someone_i [_S PRO_i to help out]
 b. someone_i [_S to help PRO_i out]

If we take the former tack, we get the reading in (38a), and we cannot fill the subject position with an overt pronoun; hence (39b), where *them* is the overt subject of the infinitival clause, eliminates this reading. If instead we map the silent pronoun into the object position of the infinitival relative clause, as in (40b), we get the reading in (38b), and now we cannot fill the object position with an overt pronoun; so (39a) is incompatible with this reading because *them* overtly occupies the object position here.

The message that (38), analysed as in (40), brings is that there can be more to the syntax of a sentence than what meets the eye or ear. Syntactic structures can be more abstract than the surface forms — in particular, they can feature silent material such as PRO. The fact that this pronoun is silent makes it impossible to see or hear where this element is sitting in the structure — and this is precisely what gives (38) its ambiguity: when PRO occupies the subject position of the infinitive, we get reading (38a); if it sits in the object position, we get (38b).

The significance of abstraction from surface reality is illustrated in a very striking way by the minimal pair of sentences in (41).

- (41) a. John is eager to please
 b. John is easy to please

On the surface, these two sentences look very much the same. They can both be characterised by the surface structural description ‘NP – *be* – A – *to* – VP’. Yet interpretively, (41a) and (41b) are quite dramatically different. In (41a), we interpret *John* as the subject of *please* (with some unspecified person being pleased), whereas in (41b) *John* is interpreted as the object of *please* (with some unspecified person doing the pleasing).

That this interpretive difference is not just a semantic accident but something that is rooted in a difference in syntax between the two sentences is perhaps particularly clear from (42) and (43). The b-examples are variants of the sentence in (41b), this time featuring *John* in the object position of *please*. For (41a), such variants are impossible, as (42a) and (43a) show.

- (42) a. *it is eager to please John
 b. it is easy to please John
 (43) a. *to please John is eager
 b. to please John is easy

Of course, the fact that we interpret *John* as the semantic object of *please* in (41b) should not make it too much of a surprise to find that it can actually occupy the object position of this verb. And conversely, the fact that *John* is not interpreted as the semantic object of *please* in (41a) makes it understandable that (42a) and (43a) do not work.

The fact that in (41b) we interpret *John* as the object of *please* should also make it possible to supply an overt subject for the infinitival clause, whereas expressing an overt subject for the infinitival clause should be impossible in (41a), where *John* is the pleaser. And indeed, as the contrast between (44a) and (44b) shows, this expectation is borne out.

- (44) a. *John is eager for *us* to please
 b. John is easy for *us* to please

Although the facts in (42)–(44) are fairly transparently rooted in the status of *John* as the object or subject of *please*, there are other syntactic differences between (41a) and (41b) for which it is not obviously the case that the semantic status of *John* is the cause. We will study one of these syntactic differences in a bit of detail: the contrast in (45).

- (45) a. John's eagerness to please
 b. *John's easiness to please

Here the fact that *John* is the subject of *please* in (45a) but its object in (45b) does not immediately help us explain the contrast. What we need is a syntax for the sentences in (41a) and (41b) that will give us an explanation for these facts.

The ill-formedness of (45b) is part of a broader generalisation, of which the example in (46b) is a further instantiation. There is nothing wrong with the sentence in (46a); but the corresponding nominalisation in (46b) does not work very well at all.

- (46) a. John appears to please her
 b. *John's appearance to please her

The parallel behaviour of (45b) and (46b) cannot be attributed to the semantic status of *John* (which is different in the two cases). Indeed, if the question of whether *John* is the semantic subject or object of *please* were key, we would expect (46b) to behave just like (45a); but quite the contrary is true.

What does seem to hold the key to the solution to the problem is that in both (41b) and (46a), *John* starts out life in syntax as a constituent of the infinitival clause embedded under the matrix predicate (*easy*, *appear*). The derivation of the sentences in (41b) and (46a) involves a syntax in which *John* originates inside the infinitival clause despite occurring on the surface in the subject position of the finite clause. The (simplified) representations in (47) are an attempt to bring this to light:

- (47) a. [_{FinClause} John is easy [_{InfClause} to please ~~John~~]]
 b. [_{FinClause} John appears [_{InfClause} ~~John~~ to please her]]

The use of strike-through in these structures is meant to indicate that although *John* is born inside the infinitival clause, it is not pronounced there because, in the course of the syntactic derivation, it moves away from its birthplace.

Movement of this sort is a syntactic operation. And syntactic operations are subject to constraints. One such constraint is that movement is easiest within the confines of a single clause, and needs a special licence when it traverses a clause boundary. The infinitival clauses to the right of the adjective *easy* and the verb *appear* in (47) make this special licence available: these clauses are transparent to movement across their boundaries. But infinitival clauses that combine with a noun (*easiness* in (45b) and *appearance* in (46b)) are what syntacticians call ISLANDS — syntactic domains which their subconstituents cannot get off of via movement. The islandhood of clauses combining with nouns is part and parcel of a much broader generalisation about nominal constructs. We will

not try to derive this generalisation here; rather, we will take it as an empirical truth, and use it to explain the ungrammaticality of the examples in (45b) and (46b). Given that the clause combining with the nouns in (45b) and (46b) is an island, the fact that these constructions are ill-formed can now be understood against the background of the representations in (48):

- (48) a. *_[FinClause] John's easiness _[InfClause=ISLAND] to please ~~John~~
 b. *_[FinClause] John's appearance _[InfClause=ISLAND] ~~John~~ to please her]]

Now that we have the contours of an explanation for the ungrammaticality of (45b) and (46b) in place, we need to return to (45a), derived from (41a). If we had *John* be born inside the infinitival clause in (41a) and then moved it into the matrix clause, we would expect (45a) to be ill-formed: after all, we have just found out that movement of *John* out of the infinitival clause in (48) is impossible. The fact that (45a) is grammatical tells us that *John* is not getting into the structural subject position of the matrix clause by movement in this case.

That *John* does not get into the subject position of the *eager* clause in (41a) via movement but is born there instead, whereas in (46a) we are dealing with movement of *John* out of the infinitival clause into the matrix clause, can be verified on the basis of examples of the type in (49):

- (49) a. there appears to be someone in the room
 b. *there is eager to be someone in the room

It is possible in (49a), with the verb *appear*, to pronounce the semantic subject of the infinitival clause *within* that clause, and to plug the structural subject position of the matrix clause up with a meaningless pronoun, *there*. But in (49b), doing this with *eager* yields a bad result. From this we conclude that the adjective *eager* itself requires a semantic subject — a role that the meaningless pronoun *there* cannot fulfil in (49b) (hence the ungrammaticality of this example) but which the referential noun phrase *John* in (41a) is perfectly equipped to serve. With *John* serving as the semantic subject of *eager* in (41a), it cannot itself be the subject of *please* as well. Even though we understand (41a) in such a way that *John* is both the one who is eager and the one who is doing the pleasing, the noun phrase *John* itself can only serve as the subject of one predicate at a time. Since *eager* needs *John* as its subject, *John* cannot directly be the subject of *please*.

So what we need for (39) is a syntax in which *John* is born as the subject of *eager*, and the role of subject of *please* is played by a silent element that has the same referent as *John* — a silent pronoun, of the same type as the one previously used in (40). Let us place (50) (the structure thus obtained for (41a)) and the one in (47b) side by side, to make it easy to compare them directly.

- (47b) _[FinClause] John appears _[InfClause] ~~John~~ to please her]]
 (50) _[FinClause] John is eager _[InfClause] PRO to please her]]

The difference between *John is eager to please her* and *John appears to please her* is that in the former, *John* is born as the semantic subject of *eager* and is linked to a silent coreferential pronoun serving as the subject of the infinitival clause, whereas in *John appears to please her*, *John* is itself born as the subject of the infinitival clause, not as that of *appear* (which has no semantic subject, as the grammaticality of (49a) had demonstrated).

The difference between (47b) and (50) is minimal on the surface: the two sentences that are output by these structures sound almost exactly alike, differing only in that (50) has *is eager* and (47b) has *appears* in the finite clause. But at a deeper level of syntactic analysis, the difference between (50) and (47b) is profound. It is responsible for the contrast between (49a) and (49b), and also for that between (45a) and (46b).

It is important to note that the key syntactic difference between (47b) and (50) lies entirely in a distinction that is not audible on the surface, when these structures are pronounced. What makes the two structures different is the nature of the element in the subject position of the infinitival clause. But this element is silent in both cases, hence not discernible when listening to the sentences as they are pronounced. The syntax has ways to bring this difference to light; but the difference as such is totally abstract on the surface. This is just one of many examples that have been accumulated over the years that show the importance of abstraction from the surface in the pursuit of syntactic analysis. The difference between the *easy to please* and *eager to please* constructions stands out as one of the clearest arguments for the role of abstraction in syntax.

As a postlude to this discussion, we will examine one more example, highlighting not only the pervasiveness of ambiguity in syntax but also the significance, at the interface between syntax and phonology, of a struck-through silent copy of a moved noun phrase.

(51) who do you want to help out?

The sentence in (51) is ambiguous: the grammatical function of *who* can either be that of subject of the infinitival clause (cf. (52a)) or that of object (as in (52b)).

(52) a. I want John to help out
 b. I want to help out John

This is very much the same type of ambiguity as the one previously encountered in (38). But there is an interesting twist to this particular case: the ambiguity can be eliminated in rapid speech by contracting *want* and *to* into the single form *wanna*:

(53) who do you wanna help out?

While (51) supports both readings in (52), the *wanna*-contraction case in (53) is unambiguous: it only has reading (52b).

What is it that makes (52a) an impossible interpretation for (53) (whereas it is perfectly available for non-contracted (51))? The answer lies in the derivational histories of the two readings for (51). In the syntax of (52b), the subject of the infinitival clause is radically silent: of course we know that someone must be playing the role of helper, but there is nothing that expresses this role overtly. In the syntax of (52a), by contrast, there *is* an overt expression of the helper: *who*. This *wh*-word starts out life as the subject of the infinitival clause, and moves to the front of the sentence via the rule of *wh*-fronting. The (simplified) representations in (54) and (55) bring this out:

(54) *who* do you want [_S ~~*who*~~ to help out]
(55) *who* do you want [_S to help ~~*who*~~ out]

In (55), *who* serves as the object of the infinitival clause and does not intervene between *want* and *to*. As a result of the fact that *want* and *to* in (55) are string-adjacent (i.e., right next to each other), the two can merge into *wanna*. But in (54), *want* and *to* are separated from one another by the struck-out silent copy of the moved *wh*-constituent, ~~*who*~~. Although this silent copy is not itself pronounced, it is present in the syntactic representation of the sentence; and because it intervenes between *want* and *to*, it prevents these two from being merged into one.

The effect that the unpronounced copy of the moved *wh*-phrase has on the contraction of *want* and *to* is one of the most dramatic illustrations of the role that silent elements play in syntax and its interface with phonology. Silent elements are abstractions from the audible and visible reality. They are theory-internal constructs. The fact that they can be ‘brought to light’, so to speak, by phenomena such as *wanna*-contraction makes the case for such abstractions: without postulating a silent copy of *who*, as in (54), we would be hard pressed trying to account for the fact that while (51) is ambiguous, (53) is not.

III.10 Universal Grammar and impossible grammars

In the syntax of the *wh*-questions in (51) and (53), the *wh*-phrase *who* is moved into clause-initial position, and a finite auxiliary (the ‘dummy’ auxiliary *do*) is inserted to its immediate right. Both *wh*-fronting and subject–auxiliary inversion are ingredients of the syntactic toolkit. They are exploited not just in the syntax of English but also in that of many other languages — related or unrelated. Thus, in Hungarian, too, when we want to question a constituent of a finite sentence, we place it to the immediate left of the finite verb: *ki ment el?* ‘who went away?’ and *kit hívtál fel?* ‘who did you call up?’ are grammatical; **ki elment?* and **kit felhívtál?* are not, because something non-verbal (the particles *el* and *fel*) intervenes between the *wh*-constituent and the finite verb.

The syntax of human natural language is regulated by a universal set of rules, together constituting Universal Grammar. Of course we all know that the world’s languages sound vastly different from one another. Many of these differences have to do with the sound system; but undeniably, there are syntactic differences between individual languages as well. So by ‘Universal Grammar’ we do not mean that all languages are the same. Rather, what we mean is that all languages draw upon a universal toolkit of rules. The rules of *wh*-fronting and subject–verb inversion are among the tools of Universal Grammar that languages have at their disposal.

The rules of syntax make reference to syntactic structures, and syntactic structures are not built randomly. There is a structure-building mechanism (which we will discuss in detail in the more specialised syntax courses in the curriculum) which regulates the construction of syntactic structures, producing such hierarchically layered representations as the ones we encountered in this segment. In such structures, it is possible to single out certain ‘privileged’ positions for the placement of certain syntactic material — for instance, the initial position in the clause (used by *wh*-fronting) and the second position in the clause (used by subject–verb inversion). There are, by contrast, no rules of Universal Grammar that make reference to, say, the fifth position in the clause, or the last or penultimate position in the clause, or the position three words in from the left or right edge of the clause. There are principled reasons why such positions never get referred to by any rule of Universal Grammar. Rules which make reference to such positions are impossible rules; grammars that include such rules are impossible grammars, fundamentally at odds with Universal Grammar.

This is important to bear in mind because it brings home the fact that syntax (surface appearances notwithstanding) is a highly organised system. Syntactic structures are built in accordance with a strict construction manual, and syntactic rules (incl. *wh*-fronting and subject–verb inversion) are defined with reference to these structures — universally. This greatly aids children acquiring their native language. The aspects of syntax that are an integral part of Universal Grammar do not have to be acquired at all: there is a blueprint of Universal Grammar in everyone of us already at birth. What the acquisition of syntax comes down to, for the most part, is the process of figuring out which of the rules of Universal Grammar are active in the language that is being acquired, under which precise circumstances these rules apply, and how the various rules interact with one another. That is by no means a trivial job — but the job would be a great deal harder if there were no Universal Grammar.