Dependency and Directionality

Marcel den Dikken

Department of English Linguistics • SEAS • Eötvös Loránd University
Research Institute for Linguistics • Hungarian Academy of Sciences

---

Chapter 2

The directionality of structure building

2 The directionality of structure building

2.1 Introduction

In mainstream generative approaches to syntactic structure building, '[b]ottom-to-top is widely regarded as the only possible order of derivation', say Phillips & Lewis (2013:24), even though 'the evidence [for bottom-up structure building] is neither extensive nor particularly well known'. They are right to point out that 'it ought to be considered as just one among several possible ways to account for certain facts about language'. The top-down alternative deserves a serious look, and a reasoned choice between the two approaches to structure building needs to be based on a careful consideration of the relevant arguments for one or the other. In this chapter, the main arguments from the extant literature regarding the directionality of structure building are placed under the microscope.

I will approach the question of how trees grow (from the bottom up or from the top down) by asking what makes them grow in the first place. The building of clausal structure is, at bottom, rooted in the projection of a predicate–argument structure, which forms the core of every syntactic expression. The standard approach in generative theory has always been to take this predicate–argument structure to be projected first, deploying the predicate and its arguments into a syntactic core that can then be grown into a full-blown sentence via the projection of functional structure facilitating the licensing of the various ingredients of the core. Section 2.2 shows that this is not the only possible starting point for syntactic structure building, and that when we examine the way clauses are made, the top-down alternative is just as good as the bottom-up standard.

Chomsky's (1965:Chapter 3) argument for bottom-up syntactic derivation was based on the cycle, a concept that is still fundamental to linguistic theorising today — indeed, in Chomsky (2001) and subsequent work, the strong cyclicity of syntactic derivations is supposed to follow inextricably from a bottom-to-top derivational model, whose conceptual rationale is that it is supposed to make syntactic derivation computationally efficient: computation proceeds in a piecemeal fashion, chunk by chunk; phases that have been completed can be 'forgotten', which lightens the memory load of syntactic computation. I will examine this argument for bottom-up syntactic derivation from a number of angles in section 2.3, showing that it is baseless.

In section 2.4, I subsequently look in detail at two extended pleas for top-down syntactic derivation. Phillips' (1996, 2003) argument is based on the syntactic reality of the 'temporary constituents' made available by it (cf. Pesetsky 1995), and exploits this to account for differences between coordination (Right Node Raising), VP-fronting and VP-ellipsis, regarding constituency conflicts and binding-theoretic exigencies. In section 2.4.1, we will examine Phillips' arguments at quite some length because they are, on their face, very compelling, and are regularly cited as such in the literature. Upon careful inspection, however, they turn out not to carry the day.

Section 2.4.2 shifts the focus from constituency to dependency — in particular, to the way in which the grammar goes about constructing filler–gap dependencies. This will set the stage for the discussion in the remainder of the book. We will examine Chesi's (2007, 2015) arguments for a bottom-up model of syntax. Chesi's discussion concentrates on three things: making intermediate steps in the formation of long-distance dependencies fall out from the system more naturally, explaining strong island effects (subject islands, adjunct islands), and accounting for parasitic gap constructions. Chesi's arguments will turn out to have significant weaknesses. But in the course of the discussion later in this book we will find that the top-down approach does indeed serve syntactic theory very well in the realm of long-distance filler–gap dependencies and the restrictions imposed on them.

2.2 What makes trees grow

Generative syntactic theories have generally placed the root of the tree at the top of the diagram, with the branches and leaves below them.[1] One of the leaves of the tree, viz., the predicate head, is supposed to provide the lexical information (argument structure) from which the lexical core of the tree can be grown. Let us look at a simple clause like (1).

(1) she kissed him

Here, the verbal root *kiss* harbours lexical information (a 'subcategorisation frame') that tells us that it takes an internal argument (ARG_{int}), which allows the root to project a branching phrase, 'VP', including besides the verb also a position for the object:

(2) $[_{\text{VP}}\ \text{V}\ \text{ARG}_{int}]]$

This VP in turn serves as the complement of a head *v*, which, for the predicate at hand, accommodates the external argument (ARG_{ext}) in its specifier. This gets us from (2) to (3).

(3) $[_{v\text{P}}\ \text{ARG}_{ext}\ [_{v\text{P}}\ v\ [_{\text{VP}}\ \text{V}\ \text{ARG}_{int}]]]$

Our toy example in (1) also contains tense information, assumed to be brought in by an additional head, T. This head takes the *v*P as its complement, and in so doing projects to TP, as in (4).

(4) $[_{\text{TP}}\ \text{T}\ [_{v\text{P}}\ \text{ARG}_{ext}\ [_{v\text{P}}\ v\ [_{\text{VP}}\ \text{V}\ \text{ARG}_{int}]]]]$

The derivation could end here, in principle, for the particular example in (1). The tense morphology (*-ed*) is now standardly assumed not to be housed under T but to be directly attached to the lexical verb (the 'lexicalist hypothesis'). The structural relation called 'Agree' (the successor of 'government' in earlier principles-and-parameters work) allows *v* to license the internal argument, and enables T to license both the external argument and the *v*–V complex. With the subject in Spec*v*P and the English lexical verb no higher than *v*, (4) will deliver the appropriate surface word order for the sentence in (1).

But it is widely assumed that English modal auxiliaries are exponents of T, base-merged there. If this is correct,[2] we need the subject of (5) to be spelled out in a position higher than Spec*v*P.

1 Recall fn. 1 in chapter 1 for some relevant discussion.

2 Though thoroughly mainstream, the argument that English modals are merged and spelled out in T has never been airtight. The key property of English modals that any theory worth its salt needs to account for is the fact that they are uniformly finite and must precede the negation particle (*not* or *-n't*) whenever this is present. The finite-only property of modals is standardly assumed to follow from their base-generation in T. But T itself has both finite and non-finite versions; so to capture the finite-only nature of English modals, one would have to assume that they can only be merged under *finite* T. This robs the base-merger of modals in T of its explanatory value: we could just as well assume that English modals are Vs that, because of their lexical specification as [+finite], can only engage in Agree relations with a [+finite] T. The fact that modals must precede negation could still be a reason to think that modals are physically in T (though this will depend, of course, on one's treatment of *not* and *-n't*). But their finiteness does not compel us to merge them in T any more than the finiteness of, say, *has* does. That *can* does not alternate with **to can* or **canning* while *has* does alternate with *to have* and *having* is a lexical quirk attributable to the English modals' uniform [+finite] specification, which does not require or fall out from merger in T.

(5) she will kiss him

If, in addition, we assume that the subject is in the same position in the tree regardless of whether there is a modal present or not, then even in (1) *she* must be in SpecTP at the point at which the syntactic structure is handed over to the interpretive components. The uniformity of subject placement in SpecTP in English and similar languages is customarily blamed on a lexical quirk that their (finite) T has: the 'EPP property'. So by this logic, (4) is not the final product of the derivation of TP: the subject must be placed in SpecTP prior to Spell-Out.

(6) $[_{TP}\ ARG_{ext}\ [_{TP}\ T\ [_{vP}\ \sout{ARG_{ext}}\ [_{vP}\ v\ [_{VP}\ V\ ARG_{int}]]]]]$

In the 1980s, in early principles-and-parameters theory (often called 'Government-and-Binding theory' or 'GB'), the element in Spec*v*P marked here as '$\sout{ARG_{ext}}$' was called a trace — a nomenclature which took 'movement' (which had always been supposed to be a metaphor) very literally as a displacement process that leaves a trace of the moved category behind. The (historically older) alternative is to assume that there are multiple tokens of the displaced constituent present in the syntactic representation. The two tokens of ARG_{ext} in (6) are customarily called 'copies' — whence the name 'copy theory'. Thus, (6) contains two copies of the external argument, only one of which is spelled out (viz., the one in SpecTP). The representation in (6) reflects this: '$\sout{ARG_{ext}}$' is a 'silent copy' of the external argument; the notational device of strike-through marks the fact that this occurrence of the external argument remains unpronounced in the physical product.

It is important to reiterate from chapter 1 that (6) is neutral with respect to the question of how placement of the external argument in SpecTP comes about. All that (6) says is that there are two tokens (or 'copies') of the external argument. The copy theory is not itself a theory of *movement*. As I pointed out in section 1.2, the idea that one of the copies arrives in a different position from the one it started out life in as a result of movement through the tree is extrinsic to the idea as such that there are two copies: what matters, if (6) is right, is that there are two tokens of the external argument in the tree, and that the one in SpecTP is the one that is pronounced. The copy theory does not make any intrinsic claims about the direction in which syntactic structures are built — 'from the bottom up' (by which we mean, for the particular case of our simple example in (1), from the predicate head and its lexical properties up to the clause) or 'from the top down'.

As it turns out, syntactic structures can be built efficaciously in either direction when we take two important regularities about such structures into account: the notions of endocentricity and extended projection.

2.2.1 Endocentricity and extended projection

At the bottom of every structure we will find a predicate head (for clauses, typically a verb). For a bottom-up approach, this means that there is always an obvious place to start building the tree. The top-down alternative has to start at the root and find its way down to the predicate. A problem here is that there is no way to predict *a priori* what the root node will be. It could be a TP. It could be larger — for instance, if specific information-structural layers of structure are activated (projections for topics and foci), or if we are dealing with a question. It could also be smaller than TP: '*Mad Magazine* sentences' (Akmajian 1984) such as (7) are root utterances whose top node may very well be no larger than a small clause, lacking T and C entirely.

(7) Barack Obama a great president?! (you must be kidding!)

But when, in nominative–accusative languages like English, the initial constituent is an explicitly nominative noun phrase (*she* in our toy example in (1)), the root node is almost certainly a finite TP. So at least for sentences such as (1) there is little doubt as to the nature of the root node.

With the nature of the root node determined, we can subsequently rely on important regularities about the internal structure of syntactic projections to predict a significant amount of the rest of the structure. We know for every utterance that there must a predicate in it: that is a given. We also know that for every syntactic projection, there must be a head: this is the age-old notion of endocentricity. A third constant is that all functional structure sits on top of a lexical core, and that this functional structure belongs in some sense to the lexical category at its core. This was codified by Grimshaw (1991) under the name of 'extended projection':

(8) an *extended projection* is the maximal projection of a lexical category plus all the functional projections belonging to that lexical category

We can think of extended projection in terms of templates or maps (as in the 'cartographic' approach) or as the elementary trees ('treelets') of Tree Adjoining Grammar (Frank 2002 and references there). We can also think of extended projection as the construction of functional 'shells' on top of the lexical core, in the bottom-up structure-building process. A third way of thinking of extended projection, one which fits in with the top-down approach, is as a predictor of structure below the functional root node. If our root node must be a TP — for instance, as in (1), because there is a nominative subject — we can make our way down to a predicate that assigns a θ-role to that subject.[3] If the internal constitution of extended projections is predictable, we can find a path from the functional root node all the way down to the predicate head.

So in our simple toy example in (1), on a top-down approach, we encounter *she*, register that it is necessarily argumental (animate personal pronouns are never predicative) and explicitly nominative, and arrive immediately at the conclusion that the root node dominating *she* as the subject is TP. once this piece of knowledge is in place, endocentricity and extended projection direct us down the tree, first to the head of TP (T), then to its complement (*v*P), then to the complement's head (*v*), to that head's complement (VP), and down to V. In the process, we are led to the discovery of a θ-position in which a silent copy of the subject can be placed (Spec*v*P). And once we are down to the verbal head, its argument-structural properties will allow us to accommodate an object for it in the same way they do in the bottom-up approach; so *him* can be fit in as well.

(9) a. *she* = nominative, argument
→ the subject of a finite clause
→ finite clauses are minimally TPs
b. $[_{TP}$ *she*
→ TP is a projection of a head (*endocentricity*, *projection*)
→ this projection accommodates not just a specifier (*she*) but also a complement
→ the complement must be a predication structure that can check T's finiteness features (tense and φ): *v*P (*extended projection*)

3 That predicate will usually be verbal. If we take the copula to be a verb, the qualifier 'usually' can be dropped for languages which lack verbless finite clauses (with 'verbless' in the sense of lacking an element expressing finiteness). If the copula is not a verbal root, we still know what to look out for once we have a TP: there is always a predicate in every TP.

c. $[_{TP}$ *she* $[_{TP}$ T $[_{vP}$
→ *v*P is a projection of a head (*endocentricity*, *projection*)
→ this projection accommodates a specifier and a complement
→ the specifier provides a θ-position for a silent copy of the subject
→ the complement of *v* must be VP (*extended projection*)
d. $[_{TP}$ *she* $[_{TP}$ T $[_{vP}$ ~~*she*~~ $[_{vP}$ *v* $[_{VP}$
→ VP is a projection of a head (*endocentricity*, *projection*)
→ the V-head is active *kissed*, whose argument structure includes an object
→ the internal structure of VP accommodates the object
e. $[_{TP}$ *she* $[_{TP}$ T $[_{vP}$ ~~*she*~~ $[_{vP}$ *v* $[_{VP}$ V=*kissed him*]]]]]

2.2.2 Guiding the guessing: The numeration

Matters are often not as simple as in this toy example. Imagine, for instance, that our lexical verb had been *saw* instead of *kissed*. We know that *saw* is transitive, hence takes a complement; but what we cannot predict is the nature of its complement: while *kissed* will usually take a nominal complement (one does not kiss propositions, to paraphrase Williams' 1983 famous dictum),[4] for *saw* we can get a nominal complement, as in (10a), or various kinds of (small) clausal complements, with (10b–e) all featuring *him* in immediately postverbal position.

(10) a. she saw him
b. she saw him out
c. she saw him to be a good person
d. she saw him cross to the other side
e. she saw him crossing to the other side
f. she saw that he was crossing to the other side

It is impossible to guess the nature of the complement of *saw* right in a top-down approach.

But it is equally impossible to guess everything right on a bottom-up approach. For instance, imagine that at some point in the bottom-up structure-building process we had arrived at the structure for *him crossing to the other side*. This structure could serve as the complement of a verb like *saw*, as in (10e) — and if it does, there will be a licenser for the accusative case feature of *him*, which makes it a reasonable guess that *him crossing to the other side* will indeed be a constituent in the complement of a verb like *saw*. But the node dominating *him crossing to the other side* could also serve as the complement of a preposition, either as in (10e′) or in a *with*-absolute such as (10e′′), or as the postcopular constituent in a pseudocleft like (10e′′′), or as the subject of a matrix clause (as in (10e′′′′)), or even as the root, as a *Mad Magazine*-type sentence (see (10e′′′′′)). The possibilities are numerous — and not obviously less numerous in a bottom-up approach than in a top-down one.

4 This may actually be a simplification. Resultative constructions such as (i) may involve something other than *her* as the complement of the verb *kissed*: on an analysis of resultatives along the lines of Hoekstra (1988) and much subsequent work, what serves as the verb's complement here is the small clause [*her awake*].

(i) she kissed him awake

(10) e. she saw [him crossing to the other side]
e′. she was talking about [him crossing to the other side]
e′′. with [him crossing to the other side], we are left with few allies
e′′′. what she is worried about is [him crossing to the other side]
e′′′′. [him crossing to the other side] would be unthinkable
e′′′′′. [him crossing to the other side]?! (unthinkable!)

A bottom-up structure-building approach can take the guessing out of the game by working with a lexical array or numeration: if the syntactic derivation proceeds on the basis of an array of lexical elements, it will be clear what to do once [*him crossing to the other side*] has been completed. But of course working with a numeration is not just possible in a bottom-up approach: a top-down syntax could likewise be equipped with this. If the lexical array includes, besides *saw*, also the verb *crossing*, then we will no longer be tempted to map *him* into the object position of the verb *saw*, and will instead postulate a verbal extended projection in the complement of *saw*. If numerations are the solution for indeterminacy in the bottom-up approach, they will serve this purpose perfectly in any alternative top-down model as well. So the top-down approach is not intrinsically worse off than the bottom-up approach when it comes to cases of indeterminacy. Arguments one way or the other cannot be based on issues arising in connection with the syntactic patterns that verbs like *see* or gerunds such as *him crossing the street* are compatible with.

2.2.3 Top-down versus left-to-right: A brief note on linear order

One might also think that OV languages are at a considerable disadvantage in a top-down approach. How could we accommodate an object in the structure if we have not come across a verb yet in the linear string?

Here we should bear in mind that a top-down approach is not exactly the same as a left-to-right linear approach. On a strictly linear approach, with little or no reference to phrase structure, it is indeed very difficult to handle OV languages. But on a top-down approach, which works with hierarchical structures, the presence of a nominative subject automatically predicts, via extended projection, an entire clausal skeleton leading down to the predicate head. Since a predicate head is securely predicted, encountering an object before the predicate head comes in leads to no particular accommodation problem because the core predication structure has already been put in place by the time the object comes along. (For more discussion of linearisation in the verb phrase, I refer the reader to section 2.3.1.5.)

2.2.4 On the internal organisation of functional structures

Whether one takes a bottom-up or top-down approach, what makes trees grow in syntax is the drive to express *grammatically licensed predicate–argument structures*. Predication structures are the cornerstones of utterances. But adult unimpaired natural language users typically do not speak in terms of 'bare' predication structures. The constituents that make a predication structure usually need to be grammatically licensed in a certain way: the arguments need case; the predicate head needs to link up with a functional category for finiteness. If we build structures from the predicate head up to the root node, we need to ensure the presence of the requisite functional categories in the higher échelons of the tree. There needs to be a system to the organisation of the functional layers of the tree: something needs to regulate this.

Chomsky (1995:Chapter 3) presented an interesting argument showing that the order in which AgrO and AgrS (the licensers for the object and the subject, resp.[5]) are merged into the structure outside the verb phrase can be derived from the locality of the movement steps that the subject and the object need to take in order to link up with their licensers. The argument, in a nutshell, runs as follows. With AgrO taking the core verbal predication (then still labelled 'VP'; the current label would be '*v*P') as its complement, as in (11a), and AgrS merged outside AgrOP, the object and the subject can each make their way to their licensing positions (SpecAgrOP and SpecAgrSP, resp.) by skipping over just a single A-position along the way, which, while deviating from the strictest notion of locality, is technically in keeping with the Minimal Link Condition if head movement (of V to AgrO, and of AgrO to AgrS) can make the position skipped and the position landed in 'equidistant' from the extraction site. But if, as in (11b), AgrS were merged immediately outside VP, with the external argument moving into its specifier position, and if further structure can only be added at the root, not inside a node that has already been constructed (the Extension Condition), then AgrO must be erected atop AgrSP. This entails that the object needs to skip over both the base position of the subject (SpecVP) and the subject's landing-site (SpecAgrSP) before it can reach its designated licensing position (SpecAgrOP). Since no head-movement operation could ever make three specifier positions equidistant (head movement forms maximally two-member chains, in which the specifier of the head of the chain, the specifier of the foot, and the complement of the foot of the chain are equidistant), movement of the object to SpecAgrOP in (11b) would violate the Minimal Link Condition, the linchpin of Chomsky's (1995:Chapter 3) theory of locality.

(11) a. $[_{\text{AgrSP}}$ ___ [AgrS $[_{\text{AgrOP}}$ ___ [AgrO $[_{\text{VP}}$ SUBJECT [V OBJECT]]]]]]

b. *$[_{\text{AgrOP}}$ ___ [AgrO $[_{\text{AgrSP}}$ ___ [AgrS $[_{\text{VP}}$ SUBJECT [V OBJECT]]]]]]

If we accept this argument (and its attributes: two Agr-projections, argument licensing via the Spec–Head relation, and the Extension Condition), we should check whether it can be reproduced in a grammar that builds structures from the root to the leaves. The active ingredient in Chomsky's explanation of the AgrSP-over-AgrOP structure in (11a) and the ill-formedness of the AgrOP-over-AgrSP alternative in (11b) is that the former but not the latter allows us to build locality-compliant connections between the arguments in SpecAgrP positions and their traces inside the lexical core. We can use exactly the same logic on a top-down approach. Nothing needs to be added to get the desired result: what a bottom-up grammar with Agr-projections for subjects and objects can do, the top-down alternative can do just as well, *ceteris paribus*.

In current minimalist syntax, subjects and objects do not necessarily need to move to specifier positions of functional categories in order to get licensed: licensing (feature valuation) is done under a government-type relation called 'Agree'. The Agree-based system no longer derives the 'split IP' structure of early minimalist syntax from the Minimal Link Condition: in fact, current minimalist syntax has abandoned this structure entirely. Today the fact that the object-licensing head (*v*) is closer to the verbal root than the subject-licensing head (T) essentially derives from the hypothesis that in order for a node to be able to merge with T, it must first be

5 In the demonstration to follow, I simplify the discussion by setting aside the projection of T (assumed to occur between the two Agr-projections), which plays no role in the argument.

categorised as verbal: it is *v* that categorises the root; so *v* must be merged before T is brought in. In other words, extended projection is at work here: T belongs to the extended projection of 'things verbal', so we must have something verbal in order to make a TP. As I already showed above, this kind of reasoning can readily be turned upside-down: if we have a T (e.g. because we have a nominative subject), we must have something verbal in its c-command domain. Once again, the top-down approach is no worse off than the standard bottom-up approach. With the help of endocentricity and extended projection, trees grow to a large extent by themselves, no matter which perspective on the direction of structure building (or structure growth) we take.

2.3 Working from the bottom up: The cycle

Why, then, do generative syntacticians working in the Chomskian paradigm *en masse* take the bottom-up approach? This is in part for historical reasons (dating back to Chomsky 1965), and in part because of the way in which the minimalist programme was laid out (in Chomsky 1995 and much subsequent work). The cycle plays a central role throughout, in various ways.

2.3.1 Clausal recursion and the cycle

Chomsky (1965) argued for bottom-up structure building based on the claim that whereas there are instances in which transformations must apply to a clause before it is embedded in a larger clause, 'there are no really convincing cases of singulary transformations that must apply to a matrix sentence before a sentence transform is embedded in it' (Chomsky 1965:133).

Confined to singulary transformations, this was probably true at the time. But once we broaden our perspective to include the generalised transformation, it is no longer true that there were no transformations that had to apply to the matrix before they could apply to an embedded domain. Early generative syntactic theory featured at least one such case. We will study this case in some detail because it turns out to be informative in many ways.

In Chomsky's (1957) approach to recursion in clausal embedding, a proform was introduced in the matrix clause before an application of the generalised transformation linked this proform to the independently constructed clause that would, after association with the proform, come to serve as the subordinate clause. For a biclausal sentence such as (12a), we would start out with two separate component parts, given in (12b), and an application of the generalised transformation would deliver the surface output by linking the subordinate clause to the proform and eliminating the latter, as in (12c). The singulary transformation responsible for the insertion of the proform (*it*) must apply to the matrix clause before the sentence transform corresponding to the embedded clause can be integrated.

(12) a. John believes that Mary is pregnant
b. [$_S$ John believes it]
[$_{S'}$ that [$_S$ Mary is pregnant]]
c. [$_S$ John believes [$_{S'}$ that [$_S$ Mary is pregnant]]]

Singulary transformations (instances of Merge) put together two pieces of structure to form a larger structure. The generalised transformation can do more than what simple Merge can do: the particular instance of the generalised transformation that integrated the subordinate clause into the structure of the matrix clause has to eliminate a piece of linguistic material, viz., the proform. In the *Syntactic Structures* model, there was, after all, just a single object position in the tree, and the proform and the subordinate clause could not *both* occupy this position.

With an eye towards restrictiveness in the theory, Chomsky (1965) abandoned the proform-based approach to clausal recursion, along with the whole notion of a generalised transformation. The proform-based analysis of clausal subordination was replaced with a direct clausal embedding analysis: the matrix verb selects the subordinate clause directly as its object, in the base component, which includes a base rule rewriting VP as V + S′.

There can be no doubt that the subordinate clause in sentences whose matrix verb is a member of the class of so-called 'bridge verbs' (Erteschik-Shir 1973) can serve as the complement of the verb, entertaining an Agree relation with the verb. But at the same time there is good reason to believe that bridge-verb complement clauses depend for their relation with the verb on the mediation of a proform, apparently very much in the way Chomsky (1957) approached the syntax of clausal complementation. We can see all of this quite directly in Hungarian.

2.3.1.1 Clausal subordination: Bridge verbs, factive verbs, and pronominal heralds

In Hungarian, the Agree relation between bridge verbs and their subordinate CPs is very clearly in evidence. Thus, in (13a), which involves a grammatical long-distance non-argument dependency between *hova* 'where.to' and a gap in the embedded clause, the matrix bridge verb must show definite inflection:

(13) a. hova akarod/*akarsz, hogy menjen?
where.to want.2SG.DEF/INDEF that go.3SG.SUBJUNC
'where do you want him to go?'

b. hova akarsz/*akarod menni?
where.to want.2SG.INDEF/DEF go.INFIN
'where do you want to go?'

That the inflectional form of the matrix verb is truly a function of the properties of the clause embedded under the matrix verb is shown by the fact that the definiteness inflection on *akar* 'want' covaries with the finiteness of the subordinate clause: Hungarian treats finite subordinate clauses as definite objects and infinitival ones as indefinites, whence the definite form *akarod* in (13a) and the indefinite form *akarsz* in (13b).

But even though Hungarian can evidently treat bridge-verb complement clauses as direct dependents of the matrix verb, it can also represent the relation between the verb and the finite object clause indirectly, via the mediation of a proform. In (14a), the non-presuppositional clausal complement is 'announced' or 'heralded', as it were, by a pronoun (*azt* 'it-ACC').

(14) a. János *azt* hiszi, hogy Mari terhes (Hungarian)
János it-ACC believes that Mari pregnant
'János believes that Mari is pregnant'

b. János hiszi, hogy Mari terhes
János believes that Mari pregnant
'János truly believes that Mari is pregnant'

The 'herald' *azt* is not strictly obligatory: (14b) is grammatical alongside (14a). But as De Cuba & Ürögdi (2009) point out, the interpretation of (14b) is not exactly the same as that of (14a): when *azt* is not present, the *hogy* 'that' clause shows a strong tendency to be interpreted presuppositionally. In this respect, (14b) is similar to so-called factive verb constructions — constructions in which the proposition expressed by the subordinate clause is presupposed to hold.

(15) a. János beismeri *azt*, hogy Mari terhes (Hungarian)
János admits/confesses it-ACC that Mari pregnant
b. János beismeri, hogy Mari terhes
János admits/confesses that Mari pregnant
both: 'János admits/confesses that Mari is pregnant'

In (15b) we see that in factive verb constructions, too, *azt* can be absent. But absence of *azt* is not a requirement in such constructions: (15a) is also perfectly grammatical.

While the Hungarian data discussed in the previous paragraph might suggest that, effectively, bridge verbs and factive verbs behave very much alike, there are significant differences between the two, manifesting themselves in various ways. The difference that is most relevant in the present context is that while *azt* alternates with *úgy* 'so' in (14a) (without a change in meaning of any kind), it does not in (15a): (16b) is ungrammatical.

(16) a. János *úgy* hiszi, hogy Mari terhes (Hungarian)
János so/thus believes that Mari pregnant
'János believes that Mari is pregnant'
b. *János beismeri *úgy*, hogy Mari terhes
János admits/confesses so/thus that Mari pregnant

This contrast points us in the direction of an important structural difference between the a–sentences in (14) and (15), regarding the way the syntax treats the pronominal element *azt*, and, concomitantly, regarding the way the grammar integrates the subordinate clause into the structure of the complex sentence. Figuring this out will tell us a lot about the way clausal subordination works in natural language, and how the grammar should operate in order to deal adequately with the syntax of clausal hypotaxis.

2.3.1.2 Object positions and presuppositionality

Let us start with the fact that, with both *hisz* 'believe' and *beismer* 'admit/confess', the subordinate clause optionally co-occurs with the pronominal element *azt*, but that with *hisz* the absence of *azt* has a presuppositionalising effect on the embedded clause. In connection with this, it is interesting to note that in Dutch and German, presuppositional complement clauses can occur in the 'middle field', to the left of the verbal cluster, whereas non-presuppositional ones never can. We see this in (17). (17a) features factive *betreuren* 'regret', for which Koster (1989) first drew attention to the fact that its complement clause can occur in non-extraposed position. In (17b), *geloven* 'believe' is used in its epistemic sense, equivalent to 'think', and we find its complement clause in extraposed position. When the complement clause of *geloven* occurs in the middle field, as in (17c), a presuppositional reading is forced upon this clause (and the matrix verb prefers to be emphatically stressed and modified by *echt* 'really'). (For more discussion of the linearisation of subordinate CPs *vis-à-vis* the matrix verb in Dutch, see the first part of section 2.3.1.5.)

(17) a. dat Jan [dat Marie zwanger is] betreurt (Dutch)
that Jan that Marie pregnant is regrets
b. dat Jan gelooft [dat Marie zwanger is]
that Jan believes that Marie pregnant is
c. dat Jan [dat Marie zwanger is] echt gelóóft
that Jan that Marie pregnant is really believes

These observations about Hungarian and Dutch taken together translate structurally into a syntactic analysis in which the verb phrase accommodates two object positions — much like the way the clause accommodates (at least) two subject positions. One object position is the familiar complement-of-V position; the other is the specifier of the projection of the verbal root (VP). This is shown in (18) (cf. Larson 1988, Hale & Keyser 1993, Barbiers 2002, Bowers 2010, Resenes & Den Dikken 2012, Den Dikken 2015, *i.a.*).[6]

(18) $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ <OBJECT> $[_{V'}$ V <OBJECT>]]]]

Whenever the occupant of SpecVP is an argument (i.e., receives a θ-role in this position), it is interpreted *presuppositionally*. How come?

We are familiar from the work of Diesing (1992) with a bipartition between the TP and the verbal domain: for quantificational subjects the nuclear scope of the quantifier can be mapped onto the verbal portion of the structure (*v*P) whereas its restriction can be mapped onto the higher portion of the clause, outside *v*P. The double vertical lines in (19) mark the boudary between the nuclear scope (to the right of the lines) and the restriction (to their left).

(19) $[_{TP}$ <SUBJECT> [T ‖ $[_{vP}$ <SUBJECT> [*v* ...]]]]

When the subject is interpreted in the Spec*v*P position in (19), it is interpreted within the domain of existential closure, and hence receives a non-presuppositional interpretation; when the subject is interpreted in SpecTP in (19), on the other hand, it necessarily receives a presuppositional interpretation.[7] There seems to be a major watershed between T and *v*P, as marked by the '‖' in (19). The Mapping Hypothesis, stated here as in (20), sums this up.

(20) *Mapping Hypothesis — subject*
- a. subject material in Spec*v*P (or lower) is mapped into the nuclear scope
- b. subject material in SpecTP is mapped into the restriction

For the 'object portion' of the structure, the facts of object shift and West-Germanic 'scrambling' suggest a similar bipartition, with a position for objects outside the domain of existential closure and one inside it. It is customary in the literature to think that the structure in (19) can be exploited for this purpose as well, with the position of shifted/scrambled objects situated outside *v*P. But there are reasons to believe that this is not correct.

6 The structure in (18) could be developed further, in such a way that the higher of the two object positions is represented as the specifier position of a functional projection between *v* and VP (e.g., AspP). Postulating such a functional projection will make the 'subject portion' and the 'object portion' of the clause structurally more similar to one another. I am very sympathetic to such a perspective. But for our purposes in this book, the simple structure in (18), where the higher object position is SpecVP, is sufficient. So to keep things simple, I will work with (18).

Krisztina Szécsényi (p.c.) asks how (18) is compatible with the Uniformity of Theta-Assignment Hypothesis (UTAH; Baker 1988). For cases in which the θ-roles assigned to the two object positions are arguably the same, a plausible response is to say that these positions facilitate the establishment of the same relation with V: in both versions of (18), the object is the argument first merged with V. UTAH is about *relations*, not absolute positions.

7 For languages like German and Dutch, the positions in which subjects are interpreted generally correspond neatly to the positions in which they are spelled out, so there are clear parallels between linear order and specificity. For English, where a sentence such as *firemen are available* is famously ambiguous between a generic and an existential reading, the subject appears to always have to be in SpecTP at Spell-out, with 'reconstruction' into Spec*v*P being possible at LF.

One is that object shift and scrambling, despite the fact that they are known to 'feed' binding relations, never manage to create a binding relation between the shifted/scrambled object and a subconstituent of the subject.[8] I illustrate this for Dutch in (21), in which it is impossible to interpret *hem* as a bound-variable pronoun linked to the universally quantified scrambled object: (21) is grammatical but supports only a referential interpretation of the pronoun *hem*; a bound-variable reading gives rise to a Weak Crossover effect.

(21) dat [telkens een andere vriend van hem] [iedere jongen] door dik en dun steunt
that every.time a different friend of him every boy through thick and thin supports
'that each time a different friend of his supports every boy through thick and thin'

The unavailability of a binding relation between *iedere jongen* and *hem* in (21) is particularly telling in light of the fact that *telkens een andere vriend van hem* 'a different friend of his each time' is saliently interpreted non-specifically. In other words, the subject is preferentially interpreted in the Spec*v*P position. The position in which the strong quantifier *iedere jongen* is interpreted must be outside the domain of existential closure. If this latter position (the 'scrambling' position) were on the edge of *v*P, above the 'base position' of the subject (Spec*v*P), it ought to be possible for the object to bind *hem*. The fact that this is impossible indicates that the 'scrambling' position cannot be above the lowest position for the subject, Spec*v*P.[9]

A second indication that a single 'specificity' position in the tree is not sufficient comes from an observation which to my knowledge originates in Ter Beek (2008:68): the fact that a 'scrambled' object can be within the nuclear scope for the subject quantifier, but can nonetheless receive a specific or generic interpretation even when the subject is clearly non-specific. One of Ter Beek's examples is reproduced in (22a) (along with the context that she provides for it); my (22b) is of the same type, and perhaps easier to judge. These examples have a bare-plural or existentially quantified subject that is the associate of *er* and therefore necessarily within the domain of existential closure, and a bare-plural object that has 'scrambled' across an adverb (*meteen* 'immediately', *hartgrondig* 'wholeheartedly') and which is interpreted generically (i.e., outside the domain of existential closure). If there were just a single bipartition of the clause into a nuclear scope and a restrictor, and this bipartition were made at the juncture of TP and *v*P, as in (19), it would be hard to imagine how (22) could be grammatical and support the particular interpretation that they do. Though probably rare, there can be no doubt that sentences of the type in (22) are grammatical. Their grammaticality suggests that a single cut between T and *v*P, for all quantificational expressions in a sentence, would be insufficient.

(22) a. [a dangerous virus has been discovered, spread by insects looking like mosquitoes]
dat er mensen muggen meteen doodslaan kan ik me dus goed voorstellen
that there people mosquitoes immediately kill can I myself therefore well imagine
'I can well imagine, therefore, that some people kill mosquitoes right away'

8 We should be careful to choose examples here that do not involve psych verbs of the *piacere*-type, for which such binding relations are known to be possible.

9 One possible way of ensuring this is to force the 'scrambled' object to 'tuck in' below Spec*v*P. But while 'tucking in' (Richards 1997) might make sense in a bottom-up derivational approach to syntax, there is no obvious rationale for it in the top-down approach that I am advocating in this work. And at any rate, 'tucking in' will not help out in the case of the second piece of data about the relationship between 'scrambling' and subject interpretation, discussed in the following paragraph of the main text.

b. ik vrees dat er altijd wel iemand generativisten hartgrondig zal verachten
I fear that there always AFF someone generativists wholeheartedly will despise
'I fear that there will always be someone who despises generativists wholeheartedly'

I propose that there is a separate bipartition into a nuclear scope and a restrictor for subjects and for objects, and that for objects, the division lies between SpecVP and the rest of the VP in (18). The continuation of the structure in (19), for transitive clauses, is thus as in (23), and on its basis we arrive at a Mapping Hypothesis for objects as in (24).

(23) $[_{TP}$ <SUBJECT> [T ‖ $[_{vP}$ <SUBJECT> [v $[_{VP}$ <OBJECT> ‖ $[_{V'}$ V <OBJECT>]]]]]

(24) *Mapping Hypothesis — object*

a. object material in the complement of V is mapped into the nuclear scope

b. object material in SpecVP is mapped into the restriction

The structure in (23) now provides us with a position for 'scrambled' objects that is at the same time *outside* the domain of existential closure for objects yet *inside* the domain of existential closer for subjects (for Ter Beek's example in (22a), this is illustrated in (25), below)) — and this is no longer a contradiction, because the nuclear scopes and restrictions for subjects and objects are computed independently of one another, via (20) and (24), respectively. In the appendix to this section (§2.3.1.5), the syntax in (23) will be put to work in a case study of verb phrase structure and linear order.

(25) $[_{TP}$ *er* [T ‖ $[_{vP}$ *mensen* [v $[_{VP}$ *muggen* ‖ $[_{V'}$ *meteen* $[_{V'}$ $[_{AP}$ *dood*] $[_{V}$ *slaan*]]]]]]]]

2.3.1.3 Back to clausal recursion

For the purposes of our current discussion, the main point of developing (23) as the structure of the transitive clause is to have at our disposal two positions for the object — one in the complement of V and the other in the specifier of VP — and to have a way of talking about the structural link between the object's occupancy of SpecVP and its presuppositional interpretation. This will now help us solve the puzzle posed by the examples in (14) and (15), viz., that in both cases the subordinate clause optionally co-occurs with the pronominal element *azt*, but that with *hisz* 'believe' in (14) the absence of *azt* has a presuppositionalising effect on the embedded clause, whereas with *beismer* 'admit/confess' in (15) the subordinate clause systematically receives a presuppositional interpretation.

In a nutshell, the proposal that I would like to advance, based on the foregoing, is the following. With factive verbs such as *beismer* 'admit/confess', SpecVP is always projected, because the complement-of-V position is occupied by a secondary predicate headed by 'FACT' — a development of the classic Kiparsky & Kiparsky (1970) approach to factives (but with 'FACT' now serving as a secondary predicate at the level of the VP, not as the head of a complex noun phrase, as in Kiparsky & Kiparsky's proposal). The occupant of SpecVP is the subject of 'FACT'. Just as in copular sentences, this subject can be either the subordinate CP itself (cf. [*that Mary is pregnant*] *is a fact*) or a proform linked to an extraposed CP outside VP (cf. *it is a fact* [*that Mary is pregnant*]).[10] Whenever it occurs in factive-verb constructions, the pronominal proform is a thematic argument of the verb. The structures in (26) summarise this.

10 (26b) is non-committal with respect to the structural position of the extraposed CP. I am generally sympathetic to an asyndetic coordination approach along the lines of Koster (2000) and De Vries (2009).

(26) a. $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ $[_{CP}$ ...] $[_{V'}$ V $[_{PRED}$ FACT]]]]]
b. $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ PROFORM$_i$=ARG $[_{V'}$ V $[_{PRED}$ FACT]]]]] $[_{CP}$ ... $]_i$

With verbs such as *hisz* 'believe', whenever *azt* is present, it once again occupies SpecVP, but this time it is not an argument but a secondary predicate for the complement clause occupying the complement-of-V position — much as in Moro's (1997) *it's that she's pregnant*, for which he treats *it* as the predicate of the *that*-clause.[11] When *azt* is not present in constructions, the subordinate clause has a choice of positions: it can either be mapped into the SpecVP position, so that it comes to behave exactly like the object clause of a factive verb (cf. (26a) and (27a′)), or it can take the complement-of-V position. The structures in (27) illustrate.

(27) a. $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ V $[_{CP}$ ...]]]]
a′. $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ $[_{CP}$ ...] $[_{V'}$ V]]]
b. $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ PROFORM=PRED $[_{V'}$ V $[_{CP}$ ...]]]]]

By Hungarian speakers for whom the presuppositional interpretation of the *hogy*-clause in (14b) is categorical, (27a′) is selected over (27a) whenever no proform is present. This can be understood from the perspective of top-down structure building: scanning in from the top/left, the first position that the object clause can be placed in is the specifier position of V. But categorical presuppositionality for (14b) is not the norm. With verbs like *hisz* 'believe', it will usually be possible to map the subordinate clause into the complement-of-V position.

Not so with factive verbs such as *beismer* 'admit/confess'. The complement-of-V position with such verbs is always occupied by an abstract secondary predicate. The proform occupying SpecVP in (26b) is an argument of that secondary predicate, and it is coindexed, moreover, with the CP in extraposed position. And in (26a), the CP itself occupies the SpecVP position. So in factive-verb constructions, the CP is always associated with the θ-role assigned to SpecVP, which ensures that it is always given a presuppositional interpretation, regardless of which of the two structures in (26) is built.

In (27b), even though the SpecVP position is occupied, we get no presuppositional reading for the subordinate clause. That is because the subordinate clause itself occupies the complement-of-V position, which is within the nuclear scope, and SpecVP is occupied by a predicate, not by the CP or a placeholder for it. Even with the predicate being assigned a presuppositional interpretation in the semantics, this does not accrue to its subject: in a copular inversion construction such as *the winner must be someone from New York*, with the predicate nominal occupying SpecTP, outside the nuclear scope, the notional subject *someone from New York* is entirely free to be interpreted non-specifically. As far as the distribution of presuppositional readings in bridge and non-bridge constructions is concerned, it thus turns out that the analyses in (26) and (27) make exactly the right predictions.

The analysis of the bridge/factive dichotomy also delivers an immediate explanation of the fact that *azt* alternates with *úgy* 'so' in bridge-verb constructions but not with factives (recall (16)). The cause of this should be plain once we realise that *úgy* 'so' is always a predicate: it has no argumental functions. Placing *úgy* in SpecVP in (26b) would cause a clash between the fact that a θ-role is assigned to this position (jointly by V and 'FACT') and the fact that *úgy* tolerates

11 Moro (1997) projects *it* in *it's that she's pregnant* as the predicate of a canonical predication structure, and has it change places with its subject (the CP) via predicate inversion (see Den Dikken 2006a). In (26b) I model the predication relation between *azt* and the CP as a 'reverse predication' or 'predicate-specifier structure', *à la* Den Dikken (2006a).

no θ-role. In (27b), by contrast, *azt* plays a predicational role; replacing it with *úgy* should be perfectly fine, and indeed it is, as we saw in (16a). We see the same alternation between *az* and *úgy* with semi-copulas like *látszik* 'seem/appear', as in (28). Here, since there is in fact no other predicate around (*látszik* is merely a copula), *az* and *úgy* play the role of *primary* predicate for the *hogy*-clause.

(28) a. *az* látszik, hogy Mari terhes (Hungarian)
it seems that Mari pregnant
b. *úgy* látszik, hogy Mari terhes
so seems that Mari pregnant
both: 'it seems/appears that Mari is pregnant'

The fact that *azt* in SpecVP in (27b) has accusative case and controls definite agreement with the matrix verb should not surprise us. Predicates in Hungarian often host case morphology. When they serve as primary predicates of a finite clause, as in (29a), they are nominative (which is morphologically invisible), under concord with the nominative subject; when they find themselves in the complement of a verb like *tart* 'consider', the case they bear is dative (see (29b)), because that is the case that the RELATOR of the secondary predication relation below *tart* happens to assign (put differently, the RELATOR = *-nak*; see Den Dikken 2006a).

(29) a. magyar vagyok, és az apám is *az* volt
Hungarian am and the father.1SG(NOM) also it(NOM) was
'I am Hungarian, and my father was, too'
b. magyar vagyok, és *annak* is tartom magam
Hungarian am and it.DAT also feel myself
'I am Hungarian, and so I feel, too'

In general, predicates in Hungarian take on the case that is available to them in their structural environment. In the structural environment in which *azt* occurs in (27b), it is the closest potential goal for *v qua* accusative case assigner. So in the environment in (27b), the pronominal predicate gets accusative case. And because the proform in (27b) is in a structural case relation with *v*, it will also control definiteness agreement with the finite verb — which hence comes out with definite inflection. When *úgy* 'so' occupies the SpecVP position instead of *azt*, definiteness agreement with it will be impossible because *úgy*, not being nominal, has no definiteness properties. And because *v* cannot establish an Agree relationship with *úgy* for definiteness, it cannot assign accusative to the proform either. So *v* skips *úgy* altogether, and targets the CP in the complement-of-V position as its Agree-goal when *úgy* is present instead of *azt*.

The analyses of clausal complementation in (26) and (27) provide a natural launching pad for an approach to so-called *wh*-scope marking constructions (a.k.a. partial *wh*-movement constructions) that finds a natural home for the '*wh*-expletive' that occurs in them. Consider a German example such as (30a), or its Hungarian equivalent in (31a) (for speaker variation on (31b), see the discussion in chapter 4 and the references cited there).

(30) a. *was* glaubt Hans, *wer* schwanger ist? (German)
what believe Hans who pregnant is
b. %*wer* glaubt Hans, dass ___ schwanger ist
who believe Hans that pregnant is
both: 'who does Hans believe is pregnant?'

(31) a. *mit* hisz János, hogy *ki* terhes? (Hungarian)
what believe.3SG.INDEF János that who pregnant
b. %*ki* hiszi János, hogy ___ terhes?
who believe.3SG.DEF János that pregnant
both: 'who does János believe is pregnant?'

In both a–examples, as in their long-distance *wh*-fronting counterparts in the b–sentences, we are dealing with root *wh*-questions. But the *wh*-operator belonging to the subordinate clause (*wer*, *ki*) is not at the left edge of the matrix clause in (30a) and (31a); the fact that we are dealing with a root *wh*-question is marked by an 'expletive' *wh*-element corresponding to English *what*. The verb agrees with this '*wh*-expletive' and assigns case to it. This is clear for Hungarian (31a) (German (30a) involves case-invariant *was*): *mit* bears the accusative case particle *-t*, and the inflectional form of *hisz* 'believe' is from the indefinite/subjective agreement paradigm, unlike what we see in (31b), where *hiszi* agrees in definiteness with the finite subordinate clause (which the grammar of Hungarian treats as definite). This property of (31a) reveals that upstairs, (31a) is the *wh*-counterpart to (14), with *mit* replacing *azt* (see Horvath 1997 for the original insight): while accusative *azt* is definite and triggers a form of the matrix verb from the definite/objective conjugation, accusative *mit* is indefinite and co-occurs with indefinite/subjective inflection. Thinking of (31a) along these lines, and bearing in mind the treatment of *azt* in (14) presented earlier in this section (recall (27b)), we immediately procure an analysis of the *wh*-scope marking construction that finds a home for the '*wh*-expletive': it originates in the SpecVP position, with the subordinate clause occupying the complement-of-V position.[12]

This also affords us an explanatory perspective on the fact that long-distance *wh*-dependencies of the b–type never allow the proform to appear in the matrix clause, regardless of whether it is [–WH] *azt* or [+WH] *mit* and (in the latter case) regardless of where the proform is placed *vis-à-vis* the animate *wh*-operator *ki*:

(32) a. **ki* hiszi <u>azt</u> hogy ___ terhes? (Hungarian)
who believe.3SG.DEF it that pregnant
b. **ki* <u>*mit*</u> hisz hogy ___ terhes?
who what believe.3SG.INDEF that pregnant
b′. *<u>*mit*</u> *ki* hisz hogy ___ terhes?
what who believe.3SG.INDEF that pregnant

The uniform ungrammaticality of the examples in (32) can now be seen to be a consequence of the fact that the *hogy*-clause is in a position where it can be an Agree-target for the verb.[13] We reach the conclusion that subordinate domains are transparent only if they are Agree-goals. This

12 A puzzle for this analysis is why *hogy(an)* 'how', which is the [+WH] counterpart to *úgy* 'so', does not occur as a 'scope marker': (i) is ungrammatical as a counterpart to (31a), despite the fact that (16a) is grammatical as an alternative to (14a). I have no definite answer to the question of why (i) fails as a scope-marking construction. It may be that scope marking generally fails when the subordinate clause is in an Agree relation with the matrix verb: in (i), *hiszi* is adorned with definite inflection, as a reflex of Agree with CP; in (31a), *hisz* has indefinite inflection.

(i) **hogy(an)* hiszi János hogy *ki* terhes? (Hungarian)
how believe.3SG.DEF János that who pregnant

13 Moreover, with *azt/mit* analysed as a predicate, extraction from the complement clause would involve a dependency between a filler and a gap inside a subject across an intervening predicate. In general, Ā-dependencies that involve a gap in the subject and traverse the subject's predicate are extremely tenuous (see Den Dikken 2006a).

conclusion, pioneered in Rackowski & Richards (2005) and traceable back to Chomsky (1986), and will be one of the central ingredients of the present work (see esp. chapter 3).

For (31b), which is grammatical for a subset of speakers, we are now led to conclude that not only is there no overt proform associated with the subordinate clause, there cannot be a 'covert' one either: the transparency of the subordinate clause to the long-distance *wh*-dependency between *ki* and the gap ('___') in the *hogy*-clause must, by the logic of the preceding discussion, indicate that the subordinate clause is itself the Agree-goal for the matrix *v*; if a nominal proform were present in SpecVP, it would be the closest Agree-goal for *v*, rendering an Agree-relation between *v* and the *hogy*-clause impossible. We see this clearly in (31a), where matrix *v* MUST Agree with *mit*. Skipping the proform and Agreeing with the *hogy*-clause is not an option: (31a) would be sharply ungrammatical with *hiszi* 'believe.3SG.DEF' instead of *hisz* '3SG.INDEF'. In order for a subordinate clause to be transparent for the establishment of a long-distance filler–gap dependency across its boundaries, the clause must be an Agree-goal to the matrix *v*. That precludes the inclusion in the matrix clause of a proform associated to the subordinate CP: when extraction from a subordinate clause takes place, CP must itself be *v*'s Agree-goal.

2.3.1.4 Conclusions and consequences concerning the cycle

The discussion in the previous subsections leads to a number of conclusions. First, it is possible for a non-presuppositional subordinate clause to be generated by itself as the verb's complement. Clausal recursion in the base component must exist, therefore. But the grammar also countenances the possibility of base-generating a proform in the higher of the two object positions (SpecVP) and associating the subordinate clause to this proform, in the spirit of Chomsky (1957): this is what happens in (26b), with factive matrix verbs. Bridge-verb constructions can mimic the structure in (26b) partially, by placing a proform in SpecVP, as in (27b). But that proform has properties that are very different from those of the occupant of SpecVP in (26b): instead of playing a thematic role, it serves as a predicate for the CP, which, in (27b) but not in (26b), occupies the complement-of-V position.

Chomsky's (1957) original proposal for the syntax of clausal subordination has now morphed into an analysis that encompasses many more facts than its predecessor and has a much wider scope, covering both bridge-verb and factive-verb constructions, making sense of the distribution and form of the proforms, accounting for extraction (see chapters 3–5), and taking care of the case and agreement facts. In Chomsky's original proposal, the clause associated to the proform is not merged into the structure of the complex sentence as an independent constituent, alongside the proform: rather, the subordinate clause is merged in via an operation of the generalised transformation, which effaces the proform and turns the embedded clause into the verb's object. The new proposal has no business with the generalised transformation: the proform, whenever present, never gets replaced; the proform and the CP each occupy their own positions in the tree, with the CP to which the proform is associated sitting in the complement-of-V position in bridge-verb constructions, and in a clause-peripheral position (likely linked to the proform via an appositional relation) in factive-verb constructions with an overt object pronoun.

At the close of this section, let us return to what originally led us to investigate Chomsky's (1957) analysis of the syntax of non-presuppositional clausal complementation: the question of whether or not it is true that 'there are no really convincing cases of singulary transformations that must apply to a matrix sentence before a sentence transform is embedded in it' (Chomsky 1965:133). Reimaged in the way outlined above, the syntax of non-presuppositional clausal complementation directly impinges on this question and on the cycle.

In the structure in (27b), repeated below, the proform must be merged into the structure before the subordinate CP is merged in the complement-of-V position. This is so because the predicative proform and the verbal root must form a complex predicate that takes the CP in the complement-of-V position as its subject. Such a complex predicate can only be formed, in the structure in (27b), if this structure is built from the top down. On a bottom-up derivation, CP is first merged directly with the verb, at the V′ juncture. At this point, CP has to be interpreted as a dependent of the verb alone. Upon the subsequent arrival of the predicative proform in the SpecVP position, we could countenance a predication relation between the complex predicate 'proform + V' only by revising the conclusion, drawn at the V′ juncture, that CP is an argument of V alone. Such a revision would amount to a derivation that is not strictly cyclic. The top-down approach, by contrast, delivers the complex predicate 'proform + V' before CP is merged into the structure. Upon merger of CP in the complement-of-V position (the last position reached in the course of the top-down structure-building process), CP is interpreted right away as the argument of the complex predicate formed by the proform in SpecVP and the V-head. No revision of a conclusion drawn earlier is necessary — the derivation proceeds strictly cyclically.

(27b) $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ PROFORM=PRED $[_{V'}$ V $[_{CP}$ ...]]]]]

So if the analysis of non-presuppositional clausal complementation in (27b) is correct, Chomsky's (1965:133) statement that 'there are no really convincing cases of singulary transformations that must apply to a matrix sentence before a sentence transform is embedded in it' must be false. The merger of the predicative proform in SpecVP is an instance of the singulary transformation: External Merge. This application of Merge must precede the merger of the complement-CP (another instance of External Merge) in order for the derivation to proceed strictly cyclically. Thus, if (27b) is right, the singulary transformation called 'Merge' must apply in the matrix clause before a sentence (CP) can be embedded in it — something that a top-down approach to syntactic structure building countenances straightforwardly, but would be problematic for a bottom-up derivation.

The outcome of the discussion of clausal complementation, triggered by Chomsky's (1957) approach to clausal recursion, thus bears in an important way on the question that is at the heart of this chapter, and of the book as a whole: the directionality of structure building. The analysis of non-presuppositional clausal complementation presented in this section makes sense only in a top-down approach to structure building, which merges the proform before the associate CP is inserted. I will elaborate further on the directionality of structure building in connection with the cycle in the next section (§2.3.2), which looks at cyclicity and opacity in tandem.

Before proceeding to this section, I will elaborate in some more detail on the structure of the verb phrase emerging from the discussion of clausal subordination, to underpin it further. This is the subject of the appendix in §2.3.1.5. Readers uninterested in the structure and linearisation of the *v*–VP system can safely skip this discussion and move straight on to section 2.3.2.

2.3.1.5 Appendix: Verb phrase structure and linearisation

The structure of the verb phrase (*v*–VP), with an 'object position' in the complement of V and another in SpecVP, can accommodate word-order variation in the verbal core of the clause if we make a small adjustment to the bottom of (23), as in (33). This structure incorporates the fundamental assumption that, universally, *v* (like other functional heads) precedes its complement, while the lexical root V is located to the right of its complement.

(33) $[_{vP}\ v\ [_{VP}\ \langle\text{OBJECT}\rangle\ [_{V'}\ \langle\text{OBJECT}\rangle\ \text{V}]]]$

The syntax of the *v*–VP system in (33) is universally the same (lexical heads take their complements to their left, functional ones to their right; specifiers are systematically on the left). Variation in the placement of the object is localised partly in the syntax (is the object in the complement-of-V position or in SpecVP?) and partly in the post-syntactic PF component (is the *v*–V complex spelled out at *v* or in VP?).[14] For English the locus of spell-out for the *v*–V complex is *v*. For consistently head-final languages such as Japanese, it is always V. The order in the *v*–VP system in the Germanic OV languages is a bit of a mixed bag: the verb is spelled out at V, to the right of the object, when the object is nominal, producing OV order; but a well-known exception to the head-finality of VP in Dutch and German (and to the head-finality of VP in other OV languages that have head-initial CPs, apparently) is the placement of CP complements to V, which must follow the verb in the linear string:

(34) a. *dat hij [dat hij Marie gekust had] dacht/zei (Dutch)
that he that he Marie kissed had thought/said
b. dat hij dacht/zei [dat hij Marie gekust had]
that he thought/said that he Marie kissed had
'that he thought/said that he had kissed Marie'

This is not the reflex of a general ban on strings of the type 'C – SUBJECT – CP – V': as we saw in (17), such strings *are* in fact grammatical, but only for CPs that are base-generated in the higher object position in (33): not the complement-of-V position but SpecVP. Dutch (35) illustrates this for the CP serving as the object of the factive verbs *betreuren* 'regret' and *ontkennen* 'deny'.

(35) a. dat hij [dat hij Marie gekust had] betreurde/ontkende (Dutch)
that he that he Marie kissed had regretted/denied
b. dat hij betreurde/ontkende [dat hij Marie gekust had]
that he regretted/denied that he Marie kissed had
'that he regretted/denied that he had kissed Marie'

Biberauer, Holmberg & Roberts (2014) blame the ungrammaticality of (34a) on what they call the Final-over-Final Constraint (FoFC):

(36) *Final-over-Final Constraint*
a head-initial category cannot be the immediate structural complement of a head-final category within the same extended projection

Because FoFC makes specific reference to the complementation relation, the grammaticality of (35a), which has the bracketed clause in the specifier of VP (with the complement-of-V position occupied by a projection of the abstract noun 'FACT'; recall (26)), does not conflict with it.

14 The proposal that (33) embodies is partly similar to Haider's (2005, 2010, 2013) work, which takes all phrases to be universally right-branching but parametrises the canonical direction of licensing for lexical heads. For Haider, in the Germanic OV languages, *v* takes VP as its complement to its right but V licenses its complement from right to left. But while Haider parametrises the direction of licensing in syntax, (33) keeps the syntax entirely constant, and relegates variation regarding the placement of the object *vis-à-vis* the verb in large part to the PF component.

Biberauer *et al.* (2014) take the bracketed clause in (34) to belong to the same extended projection as the matrix verb, thereby deviating from Grimshaw (1991) and the present work. No single extended projection can feature the same sequence of functional heads twice (see Den Dikken 2010 for pertinent discussion). The embedded clause in (34) is a full CP. Its C–T–*v* sequence cannot under any circumstances be mapped into the extended projection of the matrix verb, whose projection is dominated by a C–T–*v* sequence of its own.

We will want to derive the good result that the FoFC delivered, preserving the complement/non-complement distinction, but without turning the C-head of the embedded clause into a member of the matrix verb's extended projection. We can do this with reference to matching categorial feature content rather than membership of the same extended projection, as in the condition in (37), which is a variation on a theme repeatedly invoked in the principles-and-parameters literature, including minimalism (see Hoekstra's 1984 Unlike Category Constraint, Van Riemsdijk's 1988, 1998 Unlike Feature Condition, Heck's 2010 Constraint on Direct Recursion, and Richards' 2006 Distinctness Condition):[15]

(37) *Condition on Lexicalisation of Complementation by Lexical Categories*
a lexical head X cannot be spelled out to the right of its head-initial complement YP if X's categorial feature content is equal to or a superset of that of YP
$[x\ [_{XP}\ [_{YP}\ Y\ [_{ZP}\ ...]]\ X]]$
↑ ✳ where $\{catF\}_X \supseteq \{catF\}_{YP}$

The rationale for (37) is very much the same as that of (36): uniformity of linearisation for portions of structure that are uniform in categorial feature content. The reason for using '⊇' rather than '=' will become clear shortly, in the discussion of PP placement *vis-à-vis* the verb.

With (37) in mind, let us return to (34). Here, the matrix V takes as its complement a head-initial CP, whose categorial feature content is (given that CP is the maximal extended projection of *v*–V) the same as that of its head: [–N,+V]. In this configuration, (37) universally forces the matrix verb to be spelled out at *v*, to the left of the head-initial CP. Note that CP is structurally in the verb's complement position: the postverbal placement of the CP in the linear string does not, on present assumptions, require movement of CP into a non-complement position. This is a desirable outcome: as is well known (and as Haider 2005, 2010, 2013, in particular, has stressed repeatedly), postverbal subordinate clauses are not islands for the establishment of filler–gap dependencies across their boundaries.

In (35), because the subordinate clause is not the complement of the matrix V, (37) does not apply and the locus of spell-out of the matrix V is determined by the general setting for the linearisation parameter. Dutch sets this parameter to low spell-out of the verb, inside VP. So (35a) is straightforwardly derived. It seems plausible to take the availability of (35b) (which represents a deviation from the parametric standard not required by (37)) to be a case of analogy to (34b): subordinate clauses that serve as arguments of the matrix predicate are very often mapped into the complement-of-V position, which by (37) forces V–CP linearisation, with the *v*–V complex spelled out at *v*; by following the strategy of spelling the verb out at *v*, factive constructions assimilate themselves to non-factives, which achieves overall uniformity in the linearisation of constructions in which the matrix VP contains an argumental CP.

15 Note that (37) is not incompatible with the assumption that roots are acategorial: in the top-down approach to syntactic structure building pursued in this work, at the point at which the decision of where to spell out X is taken, X will have been provided with categorial feature content through its association with *x*, the 'categoriser'.

By (37), the *v*–V complex must be linearised to the left of a head-initial CP that serves as its complement in Dutch and other non-strictly head-final OV languages. But *v*–V otherwise generally follows its complement in these languages, even when this complement is head-initial: (37) requires that the verb be linearised to the left of its complement only when the matrix *v*–V complex has the same categorial feature content as the complement or a superset thereof. But head-initial PPs that are standardly treated as complements of V are a striking exception: their placement *vis-à-vis* the selecting verb is often, though by no means always, rather flexible.

(38) a. dat ik [$_{PP}$ aan de problemen] dacht (Dutch)
that I of the problems thought
b. dat ik dacht [$_{PP}$ aan de problemen]
that I thought of the problems
'that I thought of the problems'

It is often said in the literature that complement-PPs such as *aan de problemen* 'of the problems' in (38) can form a particularly tight-knit unit together with the verb — a kind of 'phrasal verb' *denken aan* can be formed, in one way or another (see already Hornstein & Weinberg 1981 for the roots of this idea in the generative literature). We can think of this as a case of categorial underspecification of the P in question: while the category P is standardly characterised as [–N,–V], in 'phrasal verb' or 'restructuring' contexts such as (38) P can be specified for just the categorial feature [N] (i.e., as [–N]). When P is specified as [–N] alone, its feature content is a proper subset of the categorial feature content of V (i.e., [–N,+V]). The condition in (37) is formulated in such a way that it forces structures in which a lexical head is a featural superset of its head-initial complement to be linearised with the matrix head to the left of the complement. When the P-head of the verb's complement is specified just as [–N], therefore, (37) delivers a V–PP order as the output: 'PP-over-V'. But whenever the P-head is fully specified, as [–N,–V], the matrix verb does not match or subsume the categorial feature content of its PP-complement. In such circumstances, the verb will be linearised in accordance with the general parametric setting for Dutch — that is, it is spelled out at V, to the right of PP. The variation in the placement of PP seen in (38) is thus a reflex of the richness of the categorial feature specification of the PP in the verb's complement: when PP is fully specified, it precedes the verb; when it is only specified as [–N] and thus is a category-featural subset of V, it follows the verb.[16]

16 It is important to point out that this conclusion only holds for *head-initial* PP-complements. When the PP in the verb's complement happens to be head-final (which it is systematically when P's dependent is a so-called R-pronoun; Van Riemsdijk 1978), (37) is inoperative: recall that it is formulated to apply only to cases in which X's complement is head-initial. So nothing requires pronunciation of a head-final PP to the right of the *v*–V complex in languages like Dutch, not even when P is featurally underspecified. This is accurate: even relatively heavy R-word PPs such as quadrisyllabic *overal aan* 'everywhere of, i.e., of everything' do not naturally occur to the right of V in embedded clauses:

(i) ?*dat hij had gedacht overal aan *cf.* dat hij had gedacht aan alles (Dutch)
that he had thought everywhere of that he had thought of everything
'that he had thought of everything'

Because (37) cannot force placement of the PP to the right of the verb when PP is head-final, the general parameter setting for Dutch kicks in and *v*–V is linearised to the right of the head-final complement-PP.

This also predicts, correctly, that extraction of P's immediate dependent from a complement-PP should be legitimate in Dutch only when the PP appears to the left of the verb:

In (38) we are dealing with what is standardly treated as a complement-PP, whose placement is quite free (except for the restrictions noted in fn. 16). For pre-PPs predicated of the object (raised to subject in unaccusative constructions), on the other hand, postverbal placement is often impossible. In the examples in (39a), featuring an unaccusative positional verb, and their causative counterparts in (39b), the PP *op het bed* 'on the bed' can only be placed to the left of the verb.

(39) a. dat voor de warmte mijn wollen dekens <op het bed> lagen <*op het bed>
that for the warmth my woollen blankets on the bed lay on the bed
'that to keep me warm my woollen blankets were lying on the bed'
b. dat ik voor de warmte mijn wollen dekens <op het bed> legde <??op het bed>
that I for the warmth my woollen blankets on the bed laid on the bed
'that to keep me warm I put my woollen blankets on the bed' (Dutch)

Still, it would not be right to generalise that pre-PPs predicated of the object must appear to the left of the verb. When we replace the definite noun phrase *mijn wollen dekens* 'my woollen blankets' with non-specific indefinite *wat wollen dekens* 'some woollen blankets', as in (40), postverbal placement of the PP is fine.[17] The pattern is, as far as have I been able to determine, entirely systematic.[18]

(ii) a. die problemen, daar had hij niet aan gedacht (Dutch)
those problems there had he not of thought
a′. *die problemen, daar had hij niet gedacht aan
those problems there had he not thought of
b. de problemen waar hij aan dacht
the problems where he of thought
b′. *de problemen waar hij dacht aan
the problems where he thought of

This is because the only dependents of P that can be extracted from Dutch PPs are what Van Riemsdijk (1978) called 'R-words'; and R-words precede the prepositions that select them (*eraan* 'thereof' versus **aan er*), so R-word PPs are head-final. Because of their head-finality, (37) does not apply to linearisation of R-word PPs *vis-à-vis* the selecting verb, and postverbal placement is impossible. It follows that stranded prepositions will always precede the verb in Dutch. Note that the ungrammaticality of (iia′) and (iib′) has nothing to do, on this analysis, with islandhood.

17 The presence of *voor de warmte* ensures that 'expletive' *er* is omissible in (40a). The sentence is slightly better when *er* is included, but this holds regardless of whether the locative PP is positioned pre- or postverbally.

18 I am not aware of any existing literature on the role of specificity in connection with 'PP-over-V'. What the literature *has* discussed, from a variety of different angles, is the fact that 'PP-over-V' is often facilitated by the addition of a particle to the verb (see (i)). This effect of the particle is sometimes attributed to the fact that the particle itself can be construed as a predicate, with the PP becoming an adjunct in the presence of the particle. In my own earlier work (Den Dikken 1995) I argue, however, that the PP in (ib) is still directly predicated of *het boek* 'the book', and construct a syntactic representation for (ib) that creates just the right amount of distance between the verb and the predicative PP to facilitate its 'extraposition'. I will not go into the details; they are irrelevant here.

(i) a. dat ik het boek <op de plank> zette <*op de plank> (Dutch)
that I the book on the shelf put on the shelf
'that I put the book on the shelf'
b. dat ik het boek <op de plank> *neer*zette <op de plank>
that I the book on the shelf down.put on the shelf
'that I put the book down on the shelf'

(40) a. dat voor de warmte ?(er) wat wollen dekens <op het bed> lagen <op het bed>
that for the warmth some woollen blankets on the bed lay on the bed
'that to keep me warm there were some woollen blankets lying on the bed'
b. dat ik voor de warmte wat wollen dekens <op het bed> legde <op het bed>
that I for the warmth some woollen blankets on the bed laid on the bed
'that to keep myself warm I put some woollen blankets on the bed' (Dutch)

That the specificity, not the definiteness, of the noun phrase of which PP is predicated lies beneath the acceptability of postverbal placement of the PP is shown in a particularly direct way by the facts in (41)–(42).

(41) a. dat het boek <op de plank> ligt <op de plank> (Dutch)
that the book on the shelf lies on the shelf
b. dat hij het boek <op de plank> legt <op de plank>
that he the book on the shelf lays on the shelf
(42) a. dat de nadruk <op resultaten> ligt <op resultaten> (Dutch)
that the emphasis on results lies on results
b. dat hij de nadruk <op resultaten> legt <op resultaten>
that he the emphasis on results lays on results

In both (41) and (42) the subject of the PP-predicate is definite, introduced by the definite article. Yet while *het boek* denotes a specific entity, *de nadruk* is clearly non-specific. In (41), we see the picture familiar from the earlier literature on 'PP-over-V': the PP is grammatical only in pre-verbal position. But the predicative PP in (42), with its non-specific subject of predication, can readily occur to the right of the verb.

This recalls our earlier finding that the two positions that objects can take within the VP are not equivalent semantically: the SpecVP position is outside the domain of existential closure, and hence suitable only for [+specific] objects; objects located or binding a silent copy in the complement of V, by contrast, support a non-specific interpretation.

(43) $[_{vP}\ v\ [_{VP}$ <OBJECT> $[_{V'}$ <OBJECT> V]]]
[+spec] [–spec]

To accommodate the PP as a secondary predicate of the object (raised to subject position in the unaccusative a–examples), (43) can exploit the complement-of-V position, as in (44).

(44) $[_{vP}\ v\ [_{VP}$ OBJECT $[_{V'}\ [_{PP}$ P DP] V]]]

This structure delivers two things: *(a)* a necessarily [+specific] object (because of the fact that the object is mapped into the SpecVP position, outside the domain of existential closure for objects), and *(b)* a linear order in which the PP surfaces to the left of the verb. In OV languages, V (except when located in the left periphery) can never be spelled out to the left of a noun phrase or a copy thereof. While spell-out of the verb at *v* can be forced in cases of featural harmony between V and its complement (see (37)), when the verb ([–N,+V]) has exactly the opposite feature specification as something in its domain its linearisation always follows the general parametric pattern for the language — which, in Dutch, means low spell-out, within VP.

To account for secondary predication constructions with a non-specific subject, we have to situate the entire predication structure ('small clause') within the complement of V:

(45) $[_{vP}\ v\ [_{VP}\ [_{V'}\ [_{RP}$ OBJECT $[_{R'}$ RELATOR $[_{PP}$ P DP]]] V]]]
[–spec]

In Dutch, the verb is parametrically pronounced low, in VP. I hypothesise that in (45), the lowest possible place to spell out the verb is at the RELATOR. Spelling the verb out at R is not always possible (see below); but when it happens, the verb is spelled out right in between the subject of secondary predication and the secondary predicate. Thus, (45) can derive 'PP-over-V' for PPs that serve as VP-internal predicates of a noun phrase — and it directly connects postverbal PP placement to the non-specificity of the object, which in (45) is in a position within the domain of existential closure. This, then, gets us the link between 'PP-over-V' and non-specificity.

Under what circumstances can the verb be pronounced at R, to the left of the small-clause predicate? A syntactic condition for the spell-out of the verb at R is *co-projection* of V and R, which I will talk more about in chapter 5 (*q.v.* for details). A morphological requirement for V+R co-projection is the lack of a feature conflict between V and R. This means that R should either be verbal or featurally underspecified as [–N]. In the structure in (45), R cannot be purely verbal (Dutch does not have serial verb constructions or particles that are [–N,+V]); so for co-projection to succeed, R in (45) must be [–N].[19] Now imagine that PP is itself categorially underspecified (as we know it can be: recall the discussion of (38)). Then the triad of *(a)* co-projection of V and R, *(b)* the categorial underspecification of both R and PP as [–N], and *(c)* the condition in (37) delivers postverbal placement of the predicative PP: what we have is a configuration in which the V+R co-projection complex has a head-initial complement whose feature content is a subset of *v*–V's (viz., [–N]), and which must hence be spelled out to the right of the verb. Achieving this by spelling the verb out at *v* is neither typologically desirable (because it contravenes the parametric norm for Dutch) nor observationally possible: recall from the discussion of (44) that in OV languages, *v*–V cannot be spelled out to the left of a noun phrase. Spelling the verb out at R, on the other hand, has the optimal result — an output compliant with (37) *and* with the verb pronounced within VP (as is the parametric standard for OV languages).[20]

This discussion of the placement of PP in the linear order of Dutch has thus lent support for the outlook on the structure of the verb phrase adopted in this work. The key ingredients of this perspective are *(a)* the distinction between two object positions in the structure of *v*–VP, one in SpecVP and the other in the complement of V (not necessarily the complement position of V itself; it can also be a position further embedded within the complement of V, as in (45)), each with different specificity signatures; and *(b)* the head-finality of the lexical VP, which, in conjunction with (37), gives us an account of linear order variation (both inter- and intra-linguistic). In the rest of this book, linearity will play no further role. But the general outlook on the internal structure of *v*–VP will continue to serve us very well at various points in what follows.

19 On the assumption that 'verbal particles' (Ps that form a 'phrasal verb' together with V) are underspecified as [–N], this will immediately account for the fact that the presence of a particle facilitates postverbal placement of a PP predicated of the object (see Den Dikken 1995 and fn. 18, above).

20 Of course we should not be deriving the result that whenever the PP predicate is located in an RP in V's complement, as in (45), extraposition is systematically forced: we know that 'PP-over-V' is never in fact obligatory, not even when a particle is present. We can derive this by making the eminently plausible assumption that when PP is not featurally underspecified but bears the full feature set [–N,–V], and (37) hence does not force spell-out of the verb to the left of its complement, the verb is spelled out at V, its parametric default position on OV languages.

2.3.2 Cyclicity and opacity: On the phase

In a bottom-up derivation familiar from mainstream Chomskian generative grammar, an element selected by a particular predicate is always projected at the outset (by External Merge) in its selected position. If it turns out that this element, in addition to the properties (θ-role, animacy, case, φ) that make it compatible with the selectional restrictions imposed by the predicate, also possesses a feature that cannot be satisfied in the selection position (e.g., a [WH] feature), the expression needs to establish a connection with a functional category projecting higher up the tree that has a feature matching this additional feature borne by the object. If this functional category is located outside the local domain for the element, the element must be manoeuvred, in the course of the syntactic derivation, into a position from which a local relationship with the functional category CAN be established.

The top-down perspective turns this approach on its head, merging the element in the position where its 'extra' feature is satisfied (say, SpecCP, for the [WH] feature) and tracing it back to a position local to its selector. But once again, the establishment of the filler–gap dependency must be constrained by locality considerations: a direct dependency can be established only of the filler and the gap are members of the same domain.

A central question in the syntax of filler–gap dependencies, therefore, is how the domains within which they can be established are defined. Two approaches to this have permeated the history of generative syntax — the decree-based approach and the algorithmic approach.

2.3.2.1 Opaque domains: By decree or algorithm?

The earliest approach to syntactic locality in the generative framework can be found in Chomsky (1964), couched in terms of what Ross (1967) would later call the 'A-over-A principle'. This is an algorithmic approach: for any element of type A, one can readily calculate the kinds of domains across which it can and cannot engage in a dependency with a gap. In Chomsky's (1964:931) informal formulation, the A-over-A principle

> asserts ... that if [a] phrase X of category A is embedded within a larger phrase ZXW which is also of category A, then no rule applying to the category A applies to X but only to ZXW

So in an A-over-A configuration, the grammar picks the larger of the two instances of category 'A', not the smaller one. Bresnan (1976) later develops the A-over-A into what is, with hindsight, a precursor to Rizzi's (1990) Relativised Minimality, factoring categorial information out of 'A' and stating that whenever there are two candidates in the structure for undergoing a particular grammatical operation, it should always be the larger one that is selected as the target.

Ross (1967) rejects the A-over-A approach because it is descriptively inadequate. For instance, both Chomsky's (1964) original formulation and Bresnan's (1976) development of it would, as they stand, wrongly rule out grammatical sentences such as *who would you approve of* [$_{NP}$ *my seeing* [$_{NP}$ *t*]]. In response to this, Chomsky (1968) enhances the A-over-A principle to include reference to the cycle. This helps out in the case of the example just cited: inside the gerundive noun phrase *my seeing who*, there is a cyclic node S; and within that cyclic node, the NP of *who* is the largest 'A' to which the *wh*-movement transformation could apply. Here we see the first appeal, within the context of movement transformations, to the cycle. The A-over-A principle is preserved, and the algorithmic approach to locality is continued.

Ross himself sees no virtue in pursuing the A-over-A principle further, in part because there are constraints on filler–gap dependencies that seem to have nothing to do with picking the largest target of a particular kind. Take the opacity of coordinate structures, for instance. Though one will derive the ungrammaticality of **who did you congratulate* [*Mary and* t_{wh}]*?* from the A-over-A principle if one treats the conjunction of two NPs as an NP (i.e., [$_{NP}$ [$_{NP}$ *Mary*] *and* [$_{NP}$ t_{wh}]]), the principle has nothing to say about the ill-formedness of **who did you* [$_{VP}$ [$_{VP}$ *congratulate Mary*] *and* [$_{VP}$ *commiserate* t_{wh}]]*?*, where the moved *wh*-expression is nominal whereas the coordinate structure is verbal. Ross opts instead for a list of island constraints, which Chomsky (1973) subsequently unifies (in part) under the Subjacency Condition.

Ross's island constraints and the Subjacency Condition that it gave rise to are very clear examples of the decree-based approach to syntactic locality: the inventory of 'islands' or 'bounding nodes' is settled by stipulation (for Subjacency, the bounding nodes for English are supposed to be S and NP), not computed via an algorithm. Subjacency lives on well into the principles-and-parameters era inaugurated by Chomsky (1981).

Chomsky (1986) marks a major moment in the history of syntactic locality. It presents what is to date the most explicit and detailed algorithmic approach to it: barriers are computed, for every individual dependency, on the basis of a small inventory of primitive notions ('blocking category', 'government', and 'L-marking'). Given a particular syntactic structure, the local domain for a particular element can be determined by running through the definitions. Finding out whether a trace is locally linked to its antecedent largely amounts to finding the node that dominates it, and determining whether it is L-marked (i.e., θ-governed by a lexical category) or not.

The system would have been maximally simple and straightforward if the discovery of a non-L-marked node dominating a trace that excludes the trace's antecedent had automatically doomed the filler–gap dependency. The prime reason why the *Barriers* system is so complex in its final form is that the barriers that the algorithm delivers often do not turn out to block the formation of an antecedent–trace dependency across them: more than any of its predecessors, the *Barriers* theory exploits 'escape hatches' that allow movement to proceed out of nodes that would at first blush appear to be opaque. Stopovers on the edge of a subordinate clause ('S′' in earlier work; Chomsky 1986 introduces the new label 'CP' for them) had been customary since the 1960s; now intermediate adjunction to certain blocking categories (VP, in particular; but also adjunct-PPs under particular circumstances) is added to the ways in which barrierhood can be voided. Syntactic derivations become densely successive cyclic. The empirical facts have not caught up with this theoretical development, however: unequivocal evidence for intermediate stopovers on the edges of every VP along a movement path has never materialised.[21]

21 An argument often cited in the *Barriers* era (though not in Chomsky 1986) for such intermediate adjunction was the availability of an 'upstairs' reading for *himself* (i.e., a reading in which the reflexive is linked to *John*) in ?(?)*how many pictures of himself does John wonder whether Bill bought?*. On the hypothesis that *himself* must have a local c-commanding antecedent, the *wh*-phrase *how many pictures of himself* must at some point in the derivation have been local to and c-commanded by *John*, which is the case if this *wh*-phrase makes an intermediate stopover on the edge of the matrix VP on its way up to SpecCP. Note that the matrix VP-adjoined position is the only potential intermediate landing-site for the *wh*-phrase that is local to the matrix subject: since the embedded SpecCP is occupied (by *whether*), the *wh*-phrase cannot stop over there. So if indeed *himself* in the above example is an anaphor that needs to comply with the demands of Binding Principle A, intermediate adjunction to the matrix VP (or *v*P) is necessary. But *picture*-noun phrases are well known to procure referential dependencies non-compliant with any simple version of the Binding Theory (*John thinks that pictures of himself are on sale*). BT can of course be made more complex to accommodate such cases, but it may not be a good idea to want to have BT deal with anaphors in *picture*-noun phrases. The status of examples such as the one cited above as evidence for intermediate adjunction to VP (or *v*P) thus remains uncertain. Even today, there is no clear and uncontroversial support for such intermediate adjunction. See Legate (2003) for four potential arguments, and Den Dikken (2006b) for a repartee.

Despite being cumbersome and empirically elusive, the densely successive-cyclic derivations that Chomsky (1986) inaugurated have proved an enduring legacy of the *Barriers* system: they are with us again in current 'minimalist' syntax. On the other hand, the algorithmic approach to the computation of local domains that it championed, which is conceptually attractive, has fallen by the wayside in the minimalist programme. In the more recent literature, the pendulum has largely swung back to the decree-based approach. The minimalist theory of locality, based on the phase, stipulates local spell-out domains in the form of a list (CP, *v*P, DP) very much like the way Subjacency did.[22] There is, however, an undercurrent that continues the *Barriers* line in an interesting way, represented by a proposal due to Rackowski & Richards (2005).

Rackowski & Richards argue that the establishment of an Agree relationship between the matrix *v* and the complement-CP 'opens up' the CP,[23] making it possible for *v* to subsequently probe down into the complement-CP and establish an Agree relationship with something inside it, attracting it out of CP in one fell swoop. Though this is not apparent on the surface in morphologically poor languages such as English, we see the role played by agreement in the construction of filler–gap dependencies across CP explicitly in such diverse languages as Chamorro, Tagalog, and Hungarian. Rackowski & Richards (2005) use (a particular construal of) the facts of Tagalog as their empirical test case. I will talk about Hungarian in much detail in chapter 4. So for the sake of variation, let me take the case of Chamorro (already encountered in passing above: recall (12)) for initial illustration here.

The role of agreement in long-distance filler–gap dependencies in Chamorro (an Austronesian language spoken in the Mariana Islands) jumps out particularly clearly from a comparison of the two examples in (46), taken from Chung (1998):

(46) a. hayi si Manuel hinassóso-nña chumuli' i salappi'? (Chamorro)
who Manuel WHOBJ.think.PROG-AGR WHNOM.take the money
'who does Manuel think has taken the money?'

b. guiya esti na boi i mu-na'mäguf gui' na un-li'i
he this LINK boy the WHNOM-make.happy her C WHOBJ.AGR-see
(lit.) 'this boy is the one who [that you had seen] made her happy'

As Chung (1998:250) observes, in long-distance *wh*-fronting constructions in this language, all the verbs in higher clauses are inflected for the case of the CP out of which extraction has most immediately taken place. In (46), both the downstairs verb and the upstairs one show *wh*-agreement ('WH'), as a reflex of the fact that a *wh*-filler–gap dependency is established across them. In addition, the verb in the clause in which the *wh*-expression binds a gap agrees in case with the

22 Chomsky picks CP and *v*P as phases because they are 'propositional' — complete functional complexes in the sense of Chomsky (1985). A conceptual rationale is given for why propositional categories should constitute a spell-out domain. But the rationale for identifying semantically 'whole' units as phases is undermined by the fact that, when spell-out occurs, it is not the entire phase that is shipped to the interpretive components as a unit: rather, what is spelled out is the *complement* of the phase head, which is not in any sense a semantic whole. Its interface-based rationalisation notwithstanding, in practice the identification of phases in Chomsky's work represents the decree approach. The same is true, of course, for Den Dikken's (2006a) identification of phases as predication structures. As long as the head and the edge of the phase are not handed over to the interpretive components along with the complement of the phase head, phases do not correspond to relevant entities on the semantic (or phonological) side; they may be syntactically useful entities, but their demarcation remains by and large stipulative.

23 Technically, once *v* has Agreed with CP, it may subsequently, by Richards' (1998) Principle of Minimal Compliance, ignore the complement-CP for the computation of the locality of other Agree relations.

wh-constituent in the matrix clause: 'NOM' in (33a), because *hayi* is the nominative subject of the lower clause; and 'OBJ' in (33b), as a reflex of the fact the gap in the embedded clause is in object position. But importantly, the verb of the matrix clause does not (and cannot) case-agree with the *wh*-operator: instead, it marks the case of the clause that contains the *wh*-trace. In (33a), since that clause is the matrix verb's complement, the matrix verb shows objective case agreement. And in (33b), where the clause harbouring the *wh*-trace is the nominative subject of the matrix verb, this verb shows agreement for nominative case. It is the Agree relation for case see in the Chamorro examples in (33) that opens up the subordinate clause to the establishment of a direct filler–gap dependency across its boundaries.

The key ingredient of the Rackowski & Richards (2005) approach is that Agree between a node α and a higher probe prevents opacity of α. This should apply not just to cases where α is a complement-CP: a maximally general application of Rackowski & Richards' proposal would have it that no domain Δ that serves as the goal in an Agree relation with an asymmetrically c-commanding probe could be opaque, and that, *vice versa*, every Δ that is not such an Agree-goal is impenetrable:

(47) *opaque domain*
in [α ... π ... [$_{\Delta}$ ... β ...]]], Δ is an *opaque domain* for a relation between α and β iff:
(a) Δ dominates β, and
(b) Δ ≠ a goal γ in an Agree-relation with an asymmetrically c-commanding probe π

This Agree-based approach is very much a continuation of Chomsky's (1986) notion of 'L-marking' as a preventer of inherent barrierhood. L-marking is a technically complex notion, defined as θ-government by a lexical category — a definition that is needlessly rich: government (defined in terms of minimal c-command and an opacity factor) is by and large sufficient to get the desired empirical results within the *Barriers* system.[24] 'Agree' is the successor to 'government': c-command reined in by locality. When we dress 'L-marking' down to 'government', and use the current nomenclature to refer to the government relation as an Agree relation, we end up with a generalised version of Rackowski & Richards' (2005) approach as the direct heir to the algorithmic *Barriers* theory of inherent barrierhood.[25]

24 L-marking fails, as it stands, to deliver the opacity of a CP in the complement of a noun (one of the two core cases of the Complex NP Constraint of Ross 1967): for this, Chomsky (1986) needs the entirely construction-specific stipulation that the complement of a noun is an inherent barrier (though not a blocking category). See chapter 3 for discussion. At the other end of the spectrum, a definition of an inherent barrier in terms of L-marking would make it much too hard to extract material from IP, the complement of C, and from VP, the complement of I: since C and I are not lexical categories, one would expect IP and VP to always be inherent barriers for want of L-marking. Since filler–gap dependencies that cross VP and IP on their way to SpecCP are entirely commonplace, the L-marking definition would appear to deliver the wrong results. Chomsky (1986) exploits intermediate adjunction to VP as the way to get around VP's barrierhood; but though this is technically sound, no cogent evidence for such a stopover has ever emerged (see also the discussion of intermediate adjunction to *v*P in Legate 2003 and Den Dikken 2006b). For IP, Chomsky resorts to a highly custom-made stipulation: even though IP can be a blocking category, it is never an inherent barrier. Substitution of 'government' for 'L-marking' delivers an improved result: when IP is governed by C and VP is governed by I, these categories are not inherent barriers, hence transparent, as desired. (The idea that I becomes an L-marker of VP once a verb has raised to it, which is present in *Barriers* in the analysis of NP-movement constructions, fails to support the L-marking approach because of the lack of evidence for VP's alleged inherent barrierhood in the first place.)

25 One could in principle team the approach to opaque domains based on (47) up with a decree-based inventory of 'cyclic nodes' or 'phases' — rooted in 'propositions' (Chomsky) or predication structures (Den Dikken 2006a). An Agree relation in which a node that constitutes a proposition/predication structure serves as the goal

2.3.2.2 Cyclic spell-out

This revival of an algorithmic theory of opaque domains does not combine well with the standard minimalist approach that stipulates that CP and *v*P in the clause, and DP in the extended noun phrase, are 'phases'. On the Agree approach to inherent opacity, *v*P will never be a barrier when *v* is in a local Agree relation with T, the head that selects it. CP and DP can certainly be barriers: when they occupy the structural subject position (SpecTP), they will not be c-commanded by their Agree-probe (T); when they are non-arguments, they typically will not be Agree-goals at all. But CP and DP often serve as internal arguments of verbs, and whenever they occupy the complement position of the verb, they will not be identified as opaque domains. For the syntax, this is good.

But if one wants to link the notion of 'opaque syntactic domain' to the notion of 'spell-out domain', and if one imagines the latter as an integral part of a theory that has the syntactic computation interface quickly and frequently with the interpretive components ('derivation by phase') in order to reduce computational complexity and memory load, the Agree-based approach based on Rackowski & Richards (2005) does not deliver. In a complex sentence in which multiple CPs are embedded as Agree-goals of the verb of the next-higher clause, the Agree-based algorithm identifies not a single opaque domain internal to the structure. So if opaque domains in the sense of (47) are to be the points in the derivation at which subportions of the complex sentence are 'frozen' and handed over to the interpretive components under spell-out, then in this particular kind of complex sentence there can be no sense in which the syntax liaises with the interfaces multiple times in the process of the construction of the complex sentence. If, as (47) has it, goals in an Agree relation with a c-commanding probe are systematically exempt from phasehood in the sense of 'opaque domain', and if this same sense of phasehood is to identify spell-out points, the derivation will routinely provide zero intermediate points at which the product of the syntactic Merge operations performed thus far can be handed over to the interpretive components. So if it is important for the syntax to liaise with the interfaces quickly and frequently, we will need some notion of phasehood other than (47) to tell us when to perform cyclic spell-out.

What could be the notion of 'cycle' relevant to cyclic spell-out? The answer is straightforward, regardless of whether one approaches the process of structure building from the bottom or the top. Since Grimshaw (1991), generative syntactic theory has been in the possession of a 'complete whole' that serves perfectly for the purpose of cyclic spell-out: (48).

would then exempt this node from the syntactic attributes of phasehood: i.e., it would be transparent for the establishment of filler–gap dependencies across its boundaries. In earlier (unpublished) work, I experimented with such an approach, calling it 'phase pre-emption' — the successor, in the top-down Agree-based era, to my 'phase extension' (Den Dikken 2007), *minus* the notion that phasehood shifts upwards. In the phase extension system, movement of the head of a phase to the next head up the tree achieved two things at the same time: *(a)* it eliminated the phasehood of the original predication structure, and *(b)* it handed phasehood over to the larger projection of the head that serves as the host of 'phase-extending movement'. (In Chomsky's 2015 latest work, the inverse of 'phase extension' is proposed: when C and T are in a feature-inheritance relationship, C empowers TP with phasality and is itself 'dephased'.) The upward shift of phasehood to the projection of the host of the moved head was motivated in Den Dikken (2007) primarily on the basis of the locality constraints imposed on movement in Predicate Inversion constructions. But the empirical picture here is complex (see Den Dikken 2006a) — more complex, certainly, than a pure syntax approach would be able to account for. I am no longer convinced that phasehood should be allowed to shift in the course of the syntactic derivation to nodes that are not inherently phasal, as a result of head movement or Agree. If Agree-probes are the only heads that can bring about the pre-emption of opacity as a result of instigating an Agree relation with a potential phase, and if only phase heads can be Agree-probes, the Agree approach could never make phasehood shift to a node that is not inherently phasal. I set the question aside, for lack of clear evidence.

(48) *maximal extended projection*
the maximal extended projection of a lexical root is the total set of functional projections locally dominating the lexical root which share a category specification

The largest possible extended projection of a verbal root is (the highest) CP in the functional left periphery — it is the largest possible node that has the category specification [+V,–N] (using, for convenience, the feature specifications from Chomsky's 1981 categorial matrix). For a nominal root, the largest possible extended projection is DP, the largest possible node of category [+N, –V]. Nothing requires that atop every single lexical root the syntax systematically project the largest *possible* extended projection: a clause may not need to be as large as CP; not every noun phrase is necessarily a full DP. But in any particular syntactic structure locally dominating a lexical root, the functional projection that happens to be the highest node with a particular category specification will, by (48), be declared the maximal extended projection of that lexical root. If we are to have the syntactic derivation talk regularly to the interfaces, then the maximal extended projection is a 'complete whole' ideally suited for the purpose of cyclic spell-out.

Phrased in terms of the formalism of Tree Adjoining Grammar (TAG; see esp. Frank 2002, 2006), the maximal extended projection of a lexical root is an 'elementary tree'. Elementary trees are the optimal units for piecemeal shipment to the interpretive components. In TAG, these elementary trees are in fact independent units, slotted into one another like modules, via substitution or adjoining. The elementary trees are all built autonomously, in separate workspaces. The TAG formalism is open in principle to both bottom-up and top-down approaches to the construction of the elementary trees and the larger structures composed out of them; but it does not build the entire tree fully from the bottom up or from the top down. Adjoining the auxiliary tree in (49a), with C′ as its root and foot nodes, into the elementary tree in (49b), for a full-fledged extended projection of the verb with CP as its root, delivers (49c) by creating additional structure inside an elementary subtree.

(49) a. $[_{C'}$ do $[_{TP}$ you $[_{T'}$ T $[_{vP}$ *v* $[_{VP}$ think $[_{C'}$
b. $[_{CP}$ what $[_{C'}$ (that) $[_{TP}$ he $[_{vP}$ *v* $[_{VP}$ ate *t*]]]]]
c. $[_{CP}$ what $[_{C'}$ do $[_{TP}$ you $[_{T'}$ T $[_{vP}$ *v* $[_{VP}$ think $[_{C'}$ (that) $[_{TP}$ he $[_{vP}$ *v* $[_{VP}$ ate *t*]]]]]]]]]]

This is one particular way of looking at recursion, and potentially quite an interesting one (see Frank 2002, 2006 for discussion of its virtues). But it is not one that constructs the tree in a monotonic process. I would like to explore here a tree-building strategy that is monotonic overall and fully top-down.[26]

The top-down model will always know immediately where the boundary between two extended projections lies. On the assumption (laid down in (48)) that each extended projection has a lexical root at its bottom, the derivation knows once a lexical root is reached that if anything is merged in the complement position of that root, it will define the summit of a new maximal extended projection of its own. So as soon as the lexical root anchoring the extended projection is reached, the syntax may safely conclude that the extended projection of which several members have already been included in the tree has found its way to its source. Spell-out can be triggered as soon as the grammar has verified that the lexical properties of the root have all been satisfied — including its θ-grid: I assume, in concert with Frank (2002), that (50) holds:

26 At least for complementation, or 'right-branching' structures. For specifiers ('left-branch constituents') it may not be advantageous or even possible to grow them in the process of the construction of the main projection line: it may be better to introduce these as pre-fabricated chunks. Doing so would account for the general opacity of left-branch constituents (Uriagereka 1999).

(50) all the θ-roles of a lexical root must be assigned within the root's extended projection

The thematic specifier and complement positions associated with the lexical root must hence be occupied by phrasal material before spell-out can take place. But the complement of the root does not first need to be fully completed before the structure built down to this point can be spelled out: it is sufficient for the complement position to have the appropriate phrasal occupant; the internal structure of the complement can be filled in afterwards. Concretely, then, when the top-down grammar constructs the tree for *what do you think (that) he ate?*, the matrix CP (the maximal projection of the root *think*) is spelled out as soon as a CP is postulated in the complement of *think*, as shown in (51):

(51) $[_{CP}$ what $[_{C'}$ do $[_{TP}$ you $[_{T'}$ T $[_{vP}$ *v* $[_{VP}$ think $[_{CP}$

The substructure in (51) is a complete maximal extended projection of a lexical root (*think*) within which all of the root's lexical properties have been satisfied. A complement is in place, including its label and any φ-features (for person, number, gender, or definiteness) that one might wish to attribute to it. Agreement between the matrix *v*/V complex and the complement can thus be achieved within the spell-out domain. The choice of features assigned to the complement is free in principle at this point. Though c-selection will rein in the categorial specification of the complement (English *think* is basically happy only with a CP in its complement),[27] the more microscopic properties of the complement are not necessarily predetermined: the CP in the complement of *think* could be finite, as in our toy example, or infinitival (*I didn't think to call him in time*). The assignment of features to the complement is certainly not inconsequential 'upstairs': we have already seen that in Hungarian, if a complement-CP is given the feature [+finite], it engages in an Agree-relation with the matrix verb for the feature [+definite], while a [–finite] CP controls indefinite agreement upstairs. So we will want CP in the Hungarian version of (38) to be specified for finiteness/definiteness. The assignment of a particular value for [±finite] will then predestine the complement-CP for a particular internal make-up when the internal structure of CP is built: if CP is given the feature [+finite], it will have a finite complementiser and, in turn, a finite verb; if it has the feature [–finite], it will not.

The set of spell-out domains, in a top-down model in which maximal extended projections are the loci of cyclic spell-out, includes all specifiers but only a subset of complements. Complements that are an integral part of the extended projection that is in the process of being constructed (i.e., neither the top node of a maximal extended projection nor its bottom) will never be eligible for spell-out by themselves. This conclusion converges with what the definition in (47) tells us about opaque domains in extended projection contexts. By hypothesis, all the heads in a single extended projection are in agreement with one another for at least one feature — their category specification. That is the quintessence of being in the same extended projection together. So complements that are an integral subpart of an extended projection can never be opaque domains, given (47). And we have just concluded that they cannot be subject to 'cyclic spell-out' either.

27 I acknowledge that selectional restrictions are often not specific enough to facilitate a foolproof prediction regarding the categorial identity of a selected complement (see the discussion of 's-selection' and 'c-selection' starting with Grimshaw 1979). The text example (featuring *think*, which is conveniently 'picky' regarding the categorial nature of its complement) is thus a simplification. Full determinism is impossible to achieve in the realm of complementation because of problems of categorial selection.

In neither sense of the term 'phase' (i.e., 'opaque domain in syntax' and 'point of communication between syntax and the interpretive components'), therefore, can a complement that is an integral part of an extended projection be a phase. The complement of a lexical root, on the other hand, generally heralds a new spell-out domain.[28] The top node of the complement of the lexical root belongs to the extended projection of that root (because only through its presence can the lexical root satisfy (50)). But the entire structure dominated by this top node constitutes a new spell-out domain, down to its own lexical root.

On this approach, cyclic spell-out does not track syntactic locality. The two are independent notions, not unifiable under the rubric of a 'phase'. The demarcation of a cyclic spell-out domain is arrived at on the basis of a notion of extended projection: once the *bottom* of a complete extended projection is reached, that extended projection is eligible for cyclic spell-out. The demarcation of a local domain for syntactic filler–gap dependencies, on the other hand, is a function of Agree: the *top* of an extended projection that does not serve as a goal in an Agree relation with a c-commanding probe is an opaque domain.

If the selected complement happens to be a goal in an Agree relation with a head in the matrix extended projection, as is likely, this complement is not declared a local domain, making it possible for the syntax to continue the search for a gap linked to a filler in the matrix domain. The fact that the matrix portion of the structure has already been spelled out does not affect this search. The filler does not 'disappear' from the derivation once the matrix extended projection has been spelled out; the need to associate this filler with a gap remains. But because the filler has already been spelled out along with the entire upper portion of the structure, and because spell-out allows the grammar to 'forget' about the internal affairs of the structural domain in question, it is possible in principle that when the portion of the structure that has already been spelled out gets larger and larger (i.e., the degree of embedding gets greater and greater), the grammar might 'forget' that all the way at the top of the tree there was a filler that is still waiting to be linked to a gap. The grammar might then accidentally forgo a gap altogether or postulate a gap in a position in which the filler should not be able to bind anything (which may give rise to illusions of grammaticality),[29] or perhaps postulate an 'intrusive' resumptive pronoun (in the terminology of Sells 1984). These remarks about 'forgetfulness' should be taken as preliminary; for more on the relation between the grammar and the parser, see section 2.4, below.

28 The text here says 'generally' rather than 'always' because a lexical root occasionally co-projects with the head of its complement, thereby extending the spell-out domain further down. We saw this in the discussion of PP-complementation in section 2.3.1.5. Co-projection may also be what is at work in 'restructuring' or 'clause union' constructions (see Wurmbrand 2003 and references cited there). For discussion of the notion 'co-projection', I refer the reader to chapter 5.

29 That the search for a gap can easily derail in production, with the filler seemingly being 'forgotten' altogether and no gap being provided for it, is shown by anacolutha of the type in (i). A well-known illusion of grammaticality is the fact that (iia) (an 'Escher sentence', of a type first discovered by Montalbetti 1984) initially sounds perfectly acceptable. If *more people* in the matrix clause is to identify the gap in the elliptical *than*-clause of (ia), the grammar can provide no suitable locus for this gap. We can understand the illusion of grammaticality here if the grammar 'misremembers' the comparative filler in the matrix clause — in particular, the fact that *more* is not an adverbial modifier (cf. (iib)) but instead a subconstituent of the argumental noun phrase *more people*.

(i) I was watching a programme about the Azores the other day, which I would really like to have an opportunity within the next couple of years to go on a holiday with my wife and explore the natural beauty and sample the local cuisine

(ii) a. more people visited Rome than I did
b. people visited Rome more (often) than I did

2.3.2.3 Directionality

Cyclic spell-out works fully deterministically in a top-down grammar, and involves no delays of any kind. Once an extended projection has been traced down to its lexical root and the lexical root has satisfied its selectional restrictions, spell-out of the maximal extended projection in place at this time is triggered. In a bottom-up model of structure building, on the other hand, there will be plenty of cases in which, at the point at which we reach α, a node in the extended projection of the root √, the bottom-up derivation will not be able to know with certainty (unless the grammar works with a pre-selected lexical array) whether or not α will be the maximal extended projection of √: only if α happens to be the *largest possible* extended projection of the lexical root will the bottom-up derivation be able to spell it out right upon its completion; in all other cases, the cyclic spell-out process must wait until the next head (if any) further up the tree has been merged. So if we take the maximal extended projection to be the locus of spell-out, the top-down approach has a distinct advantage over its bottom-up alternative: the top-down approach will always know the summit of the extended projection (because that is what it starts out from), and closes it as soon as the lexical root is reached and it has been verified that the lexical root has satisfied its lexical properties; the bottom-up approach, by contrast, approaches the summit from below but often will not know with certainty whether the summit has been reached until a portion of the superordinate structure is in place.

This recalls familiar problems with the spell-out of Chomsky's 'phase'. If a phase were spelled out integrally right at the point at which it is completed in the bottom-up process of structure building, we would never be able to access anything inside the phase in the continuation of the upward-bound syntactic derivation. It is customary, therefore, to build in a delay mechanism into the spell-out of phases: a phase is spelled out only at the completion of the projection of the next head up the tree, or even the next phase.

The Agree-based approach to opacity embraced in section 2.3.2.1 would, on a bottom-up approach, encounter a very similar delay problem. If determing whether a constituent α is an opaque or transparent domain involves determining whether it is in a 'government' or 'Agree' relation with a probe higher up the tree (as in the *Barriers* theory of inherent barrierhood and in Rackowski & Richards' 2005 successor thereof), then syntax will be much better served by a structure-building mechanism that creates trees from the top down than by the alternative bottom-up approach that is standard in the minimalist era.

Imagine that, in a bottom-up theory of the building of syntactic structures and the establishment of filler–gap dependencies within them, the syntax at some point has postulated a gap and is on the look-out for a local antecedent for it. When the structure-building process completes a constituent α dominating the gap, the question that arises is whether this constituent will or will not be a local domain for the gap. If local domains are defined in terms of Agree, α will be a local domain if and only if it is NOT the Agree-goal of a probe higher up the tree. So the answer to the question of whether α is or is not a local domain cannot be given instantly upon the completion of α: we will have the answer only once the next probe has been merged into the structure and we have determined whether this probe does or does not engage in an Agree relation with α.

This introduces a delay, one which undermines the determinism of the system. Some readers might interject that the bottom-up approach could encode the status of α as an Agree-goal directly on α — for instance, in the form of an uninterpretable or unvalued feature on α. Indeed, uninterpretable/unvalued features on α will identify α as something that must engage in an Agree relation with a higher probe. But it is not the case that probe-goal relations systematically affect the uninterpretable/unvalued feature on the goal: for an Agree relation involving φ-features, it

is in fact the feature bundle on the probe that is uninterpretable/unvalued; the φ-features of the goal (say, a first person pronoun, or a plural DP) are meaningful and valued. It seems unlikely, therefore, that the status of α as an Agree-goal can consistently be signalled early, in a bottom-up approach, in terms of α's feature content.

Imagine now that we built syntactic structures in the opposite direction, from the top down. Then, at the point at which we merge α, we will know right away whether it is an Agree-goal or not. By the time of α's merger, the structure will already include the superordinate portion of the tree, above α. So if that superordinate layer of the structure contains no probe whose features match those of α, we will know immediately that α is not an Agree-goal, which will render α opaque. Decisions about opacity are immediate and deterministic in this approach.

On a top-down approach in which cyclic spell-out domains are defined as maximal extended projections of a lexical root within which the lexical properties of the root are satisfied, and opaque domains are identified as nodes that do not serve as Agree-goals to a c-commanding probe, both spell-out cycles and opaque domains can be determined entirely deterministically. At the point at which a structure such as (51) is in place, we can converge immediately upon an answer to the question of whether the lexical root's complement is opaque or transparent, and the matrix extended projection can at once be targeted by cyclic spell-out and marked as complete and unalterable. This portion of the structure can then be handed over to the interpretive components, prepared for phonological and semantic interpretation, and 'forgotten'.

2.3.3 Feature inheritance and the structural subject

The Extension Condition was introduced in Chomsky (1995:Chapter 3) as the way to capture the effects of the transformational cycle of earlier models. By its very nature, this requirement (which demands that every instance of Internal or External Merge extend the root of the tree) is a linchpin of bottom-up structure building. Chomsky (1995:Chapter 4) recast the Extension Condition as a definitional property of strong features: a strong feature must be checked before its bearer is included in a larger syntactic structure that no longer has the bearer's label. In Den Dikken (2014), I show that this definition, once updated with 'EPP property' substituted for 'strong feature', would be compatible with the feature-inheritance approach to EPP-driven movement of the subject (Chomsky 2008 and subsequent work) — the idea that the subject can raise to SpecTP only after TP has merged with C, the phase-head bearer of the relevant EPP property. Though late movement of the subject to SpecTP after C has already been merged violates the Extension Condition of Chomsky (1995), that is no longer a concern once the Extension Condition is recast as a property of a feature (strength or 'EPP'). Late subject movement is not technically countercyclic: it happens fully within one single phase, and does not violate the condition imposed on the checking of strong/EPP features.

But though late subject movement does not violate the cycle, neither is it a ringing endorsement of bottom-up derivational syntax. If the subject can only raise from *v*P to SpecTP after C is present in the structure, there is a point in the derivation at which we descend back down the tree to perform an operation (NP-movement) wholly within a portion of the tree (TP) that is part of a larger structure that has a different label (CP). The derivation does not strictly speaking proceed in a purely bottom-up fashion, therefore. From a top-down perspective, the concept of downward feature inheritance, if it should turn out to be needed, is entirely unproblematic, and so, of course, is placement of the subject in SpecTP after C has been merged: that is the natural order of things.

The notional subject of the sentence is not always pronounced in SpecTP, however. In languages like English, whenever the subject is further downstream, an element is called upon to fill the structural subject position, SpecTP, in order for the 'Extended Projection Principle' (EPP) to be respected. For the elements that plug the structural subject position in case the notional subject does not (the *there* of *there*-existentials like (52a), and the *it* of sentences such as (52b), it is entirely standard in the generative literature to refer to them as 'expletives', meaningless placeholders. From a bottom-up perspective on structure building, 'expletives' make decent sense: the θ-role of the 'associate' is assigned early, and because the bearer of the θ-role ends up not moving into the athematic structural subject position, this position is plugged up with a meaningless element later in the derivation. But from a top-down perspective, it would seem awkward to merge a meaningless element in the structural subject position early in the derivation: what could motivate the insertion of an expletive in SpecTP in a top-down derivation?

(52) a. there is someone on the roof
b. it is clear that this isn't going to work

Two things should be said in answer to this question. First, it is far from obvious that *there* and *it* are truly meaningless elements in *there*-existentials and CP-extraposition constructions. Some have argued (Williams 1994, 2006, Hazout 2004) that *there* is in fact an argument, with the associate serving as its predicate; likewise, Bennis (1986) argues for the *it* of *it is clear that S* that it is an argument. Others (esp. Moro 1997) have taken *there* and *it* to be predicates, with the associate as their subject. There are a variety of *bona fide* arguments for these proforms not being meaningless, not being mere placeholders (see also Kayne 2016).[30] If *there* and *it* have meaning, the fact that they are merged in SpecTP on the top-down approach serves a meaningful purpose.

Secondly, even if *there* and *it* originated in SpecTP as truly meaningless elements, they would still serve the purely formal purpose of satisfying the 'EPP property' of T: if there is such a thing as the EPP,[31] then satisfying it is a motive for insertion of some material (whether meaningful or not) in SpecTP. In a top-down syntax, the SpecTP position is predictable: extended projection guarantees a TP in the complement of C, so once C is in place, we can predict the nature of its complement in the top-down derivation; and if for TP the EPP guarantees that it has a specifier, then once TP has been predicted, SpecTP comes along for free as well. There really is no indeterminacy in this process: merging C leads to TP; merging TP leads to SpecTP and something occupying it. If it so happens that nothing meaningful is merged in this position, a true expletive can take this spot. The top-down approach is thus not incompatible with true expletives *per se* — but it is important to note that such elements can only be postulated, in the top-down approach, in positions whose existence in the tree is independently predictable. The EPP is the only predictor of an athematic specifier. And the EPP, as formulated in Chomsky (1981) (i.e., not the 'generalised EPP' of minimalist work, which is a very different kind of creature, delivering athematic specifiers even for projections that already have a thematic one), is confined in scope to the structural subject position, SpecTP. So on a top-down approach we expect true expletives (meaningless placeholders) to be able to occur only at SpecTP.

30 In this connection, also recall the discussion of the Hungarian 'herald' pronoun *azt* in clausal subordination constructions in section 2.3.1.

31 See chapter 5 for discussion of ways of recasting the original EPP of Chomsky (1981) (i.e., the requirement that every clause have a subject).

True 'object expletives' cannot exist: there is no principle of the grammar that requires the high object position (SpecVP) to be projected or occupied; no formal considerations could possibly motivate the merger of a true expletive in an athematic position whose presence in the structure is not independently guaranteed. So all instances of proforms that have been claimed to be pleonastic and which appear in a position other than SpecTP ('*wh*-expletives' in *wh*-scope marking constructions; 'object expletives' in association with extraposed clauses) must be meaningful in one way or another. The proform *azt* in sentences such as (14a) and the upstairs *wh*-elements *was* and *mit* in (30a) and (31a) are often referred to in the literature as 'expletives'. The pronoun *azt* in factive (15a) is also commonly identified as an 'expletive' — see e.g. Postal & Pullum's (1988) discussion of 'object expletives' in English sentences such as *I (dis)like/don't mind/doubt/hate/regret/resent it that S*.

(14a) János *azt* hiszi, hogy Mari terhes (Hungarian)
János it-ACC believes that Mari pregnant
'János believes that Mari is pregnant'

(15a) János beismeri *azt*, hogy Mari terhes (Hungarian)
János admits/confesses it-ACC that Mari pregnant
'János admits/confesses that Mari is pregnant'

(30a) *was* glaubt Hans, *wer* schwanger ist? (German)
what believe Hans who pregnant is
'who does Hans believe is pregnant?'

(31a) *mit* hisz János, hogy *ki* terhes? (Hungarian)
what believe.3SG.INDEF János that who pregnant
'who does János believe is pregnant?'

But the approach to these proforms presented in section 2.3.1, above, identifies them as meaningful material: a secondary predicate in the former, with CP as its subject (see (27b)), and an argument in the latter (see (26b); cf. also Rothstein 1995).

(26b) $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ PROFORM$_i$=ARG $[_{V'}$ V $[_{PRED}$ FACT$]]]]]$ $[_{CP}$ ... $]_i$
(27b) $[_{vP}$ SUBJECT $[_{v'}$ *v* $[_{VP}$ PROFORM=PRED $[_{V'}$ V $[_{CP}$ ... $]]]]]$

Summing up this brief discussion, we have seen that feature inheritance, which is an anomaly in a strictly bottom-up derivation, fits straightforwardly into a top-down approach to syntactic structure building, and that there is no friction between this model and the possible occurrence of true expletives in the structural subject position, SpecTP, whose inclusion in the structure is fully predictable on standard assumptions.

2.3.4 Successive cyclicity

In the minimalist programme, Chomsky has sought to motivate the need for bottom-up successive-cyclic derivation in a variety of different ways. What might appear to be the most cogent one is the argument rooted in the workings of Spell-Out, the operation that 'removes LF-uninterpretable material from the syntactic object K and transfers K to the phonological component' (Chomsky 2001:5).

If this is what Spell-Out does, then it must be able to recognise as such the material that it is supposed to remove. Chomsky (2001:5) assumes that 'the uninterpretable features, and only these, enter the derivation without values, and are distinguished from interpretable features by virtue of this property'. So what makes material subject to removal under Spell-Out is its lack of a value. Values are assigned to features in an Agree relationship (basically the successor of 'government'), of which probes are the instigators. Probing features are, by hypothesis, the privilege of phase heads, so all probe–goal Agree relationships are established within a single phase. Valuation immediately makes uninterpretable features indistinguishable from interpretable ones, and therefore no longer subject to removal by Spell-Out. For the timing of Spell-Out, Chomsky says that this has the consequence that it 'must apply shortly after the uninterpretable features have been assigned values (if they have not been assigned values at this point, the derivation will crash, with uninterpretable features at the interface)', and that hence, Spell-Out 'must be strongly cyclic' (Chomsky 2001:5).

This does not follow at all; in fact, quite the opposite seems necessary if what Chomsky says is correct. Let us see why. Valuation is done by probes. Probes define phases. Spell-Out takes place at the phase. Once valued, features are no longer uninterpretable if (as Chomsky has it) 'uninterpretable' and 'unvalued' are biuniquely correlated. Spell-Out cannot recognise any 'uninterpretable' feature at the phase unless it can look back at the derivational history of the phase, discover that a particular feature started out life without a value, and determine whether or not this feature received a value from the probe heading the phase. Chomsky (2001:12) realises this but suggests that such lookback is fairly innocuous: 'The valued uninterpretable features can be detected with only limited inspection of the derivation'. But Epstein & Seely (2002) argue that such inspection of the previous derivation must be able to reach back quite a distance. In a sentence like *what do you think that Bill ate?*, the *wh*-constituent should be spelled out in the specifier position of the matrix CP. But it had its case feature valued in the *v*P of the subordinate clause. If we assume that the only uninterpretable feature of a *wh*-constituent like *what* is its case feature,[32] then Spell-Out needs to look all the way down to the *v*P of the embedded clause to find *what*'s case feature in an unvalued state and to recognise it as an uninterpretable feature. But that would mean looking deep into the bowels of the embedded CP, which is supposed to have been spelled out *in toto* by the completion of the matrix *v*P at the latest.[33]

32 This is the received view in the literature, which generally takes the [WH]-feature of *wh*-expressions to be interpretable. This is a mistake (the *morphological* feature [WH] is shared by constituents with wildly different functions and interpretations, including not just question words but also operators in relative clauses and exclamatives, and, in some languages, bare indefinites; there is no common denominator besides operatorhood that unifies these uses — and plainly 'operatorhood' is not the privileged interpretive correlate of the morphological feature [WH]). But since this mistake is pervasive and since all of the technical literature on derivation by phase and computational efficiency assumes that [WH] is interpretable, I have chosen to adopt this assumption in the main-text discussion, for the sake of argument.

33 Epstein & Seely (2002:75) suggest that Spell-Out does not apply to any particular *representation* but is part and parcel of the valuation process itself: 'we propose that Spell Out operates on all and only those formal features that appear without a value in the input to a rule being applied, but appear in the output of that rule application with the "previously" unvalued feature now valued'; 'If Spell Out can see both the input to and the output of Agree, it can see the "process" whereby an unvalued feature became valued, and then can spell out just these features, as desired.' But one wonders whether it is 'as desired' that Spell-Out affects just the features that are valued by an application of Agree: if it is literally just the features valued under a particular application of Agree that are spelled out, the derivation constantly delivers individual features to the interpretive components — not words or constituents bearing these features. In Epstein & Seely's (2002) alternative to phase-based Spell-Out, the process of 'stripping away' valued uninterpretable features is immediate and strictly derivational. This is not 'derivation by phase' but

If derivation by phase is designed to allow the syntactic computation to 'forget' earlier stages of the derivation and lighten the load for the computational system, then Chomsky's (2001) approach to Spell-Out and uninterpretability is singularly unhelpful. If, in order for Spell-Out to recognise its targets, we must retain information about a valued feature's erstwhile lack of a value over extended stretches of the syntactic derivation, there really is no benefit to the idea that the syntactic derivation proceeds phase by phase. Nor will long-distance filler–gap dependencies benefit in this system from intermediate movement via the edge of a phase. For if at the matrix CP of *what do you think that Bill ate?*, Spell-Out has access to featural information that dates back all the way to the beginning of the derivation, that information has not been 'forgotten': matrix C should have access to that information, which should enable it to attract *what* straight to its specifier, without any intermediate stop-overs being necessary along the way.

One could respond to the discussion in the previous paragraph in a number of ways. One sensible response will be to eliminate the notion of 'uninterpretable feature', and to frame the syntactic derivation entirely in terms of valuation of unvalued features. But if valuation is done by probes under Agree, and if an Agree relation between a probe and a goal that are not separated by a phase head does not require movement, the need for valuation by itself is going to give rise to very little displacement. Displacement strictly for valuation purposes could happen only when a probe and its goal are separated by a phase head that does not value the relevant feature: the goal would then need to be manoeuvred onto the edge of the lower phase in order for a feature-valuation relationship of Agree to be locally establishable between the higher probe and its prospective goal. But note that in a bottom-up derivation the goal has no way of knowing ahead of time whether a probe with a matching feature will merge outside the phase that the goal is in.

The derivation can proceed deterministically if the local phase head itself engages in a feature-valuation relation with the goal. But that local probe–goal relation (a case of Agree) does not by itself lead to displacement of the goal. Displacement is strictly the province of an extraneous factor: the 'generalised EPP', the requirement that says of a head that it needs a specifier. We are free to give the 'EPP property' to any probe, and therefore we are free in principle to give it to all probes along the path of long-distance *wh*-movement. If we choose to so do, displacement will proceed via a succession of local steps — so-called successive-cyclic movement.

In such an outlook on long-distance movement, the succession of local steps is entirely the result of the assignment of the 'EPP property' to all the probes along the movement path. Insofar as probes are by definition heads of phases, movement will indeed proceed from phase edge to phase edge. But the phase-to-phase nature of the movement is secondary: the primary active ingredient in this is the assignment of 'EPP' to a variety of heads along the way. If 'EPP' could be randomly assigned to (functional) heads, not necessarily those that are by hypothesis the heads of phases, there would be no connection between stepwise movement and phases.[34]

it is fully bottom-up and piecemeal — 'derivation by application of Agree' is probably the best way to characterise this approach. I will not discuss it further here because the details of making Spell-Out an integral part of the valuation operation have, to my knowledge, never been sufficiently developed.

34 The functional head over which the original Extended Projection Principle of Chomsky (1981) had jurisdiction was I, now usually called T. This head is not recognised as a(n inherent) phase head in Chomsky's work since the introduction of the notion of a phase. But T IS held responsible for the displacement of something (the thematic subject, or some other constituent, even an expletive) to the structural subject position, SpecTP. So T can be endowed with the 'EPP property'. But if T is not a phase head, and only phase heads are to be inherently equipped with 'EPP', how does T end up with this property? We saw in the previous subsection that a 'feature inheritance' relation between C and T can take care of this. But once 'EPP' is no longer, on the surface, a property of phase heads only, the case for tying this property exclusively to phase heads underlyingly is weakened.

Note also that if the notion 'EPP position' is considered equivalent to 'Spell-Out position', there could be no movement via intermediate steps at all. Movement to an 'EPP position' would then spell the moved constituent out right there, and prevent it from engaging in further movement operations. All things considered, therefore, the 'generalised EPP' puts successive-cyclic movement on less than solid grounds.

In his most recent work, Chomsky (2013, 2015) has suggested that the successive cyclicity of movement can be derived from the labelling algorithm — in particular, from the difficulty arising in the labelling of [XP YP] structures, tying two non-minimal projections together under a single node. Which label should we assign to [XP YP] structures? That of X or that of Y? Chomsky has suggested that the computational system cannot figure this out unless either *(a)* X and Y are in some sense the same (i.e., they share one or more features under agreement, and it is these shared features that label the node) or *(b)* either XP or YP moves on.

It is *(b)* that is supposed to derive successive-cyclic movement. Suppose that at some point in the derivation, we arrive at a structure in which a phrase 'XP' occupies the specifier position of a phrase 'YP'. The resulting [XP YP] structure cannot be labelled unless *(a)* or *(b)* holds. Suppose further that XP and YP do not share any features on the basis of which the node could be labelled — i.e., *(a)* is not in effect. That then leaves *(b)*: one of the two phrases must move on. For the particular cases in which 'XP' is a *wh*-phrase and 'YP' is a subordinate CP or *v*P, this delivers the need for one of the two terms in the [XP YP] structure to move on; and given that (at least in a language such as English) YP is stationary in these environments, it follows that it is the *wh*-phrase that has to move on.

A few things need to be noted about this proposal. The first is that, because it implicates the labelling problem incurred by [XP YP] structures in general, it does not specifically make reference to the phase. Like its predecessor, the labelling-based approach to successive cyclicity does not *per se* derive the presumed 'phase-edge to phase-edge' nature of successive-cyclic movement. [XP YP] configurations are not the privilege of phases. Perhaps [XP YP] structures that are created via Internal Merge ARE confined to phases — but then we first need to derive this before we can hope to derive the 'phase-edge to phase-edge' nature of long-distance filler–gap dependencies. The labelling approach fails to do so. Though it does potentially derive the need for a *wh*-phrase to move on from an intermediate landing-site after it got there, labelling as such does nothing to derive the movement of the *wh*-phrase to the stop-over position in the first place. Labelling might derive what happens *after* the intermediate movement step has been taken; but it does not derive this movement step itself. If this movement step is to be triggered, the only thing in the standard toolkit that can do that is the 'EPP', as I showed previously. And if 'EPP' is not an autonomous feature of a probe but instead a property of a feature of a probe, then C and *v* must have a particular feature to which 'EPP' is associated, and the moved *wh*-constituent must have the matching feature. If so, C/*v* and the *wh*-phrase will engage in an agreement relation for the feature in question, which should label the [XP YP] structure on the basis of *(a)*, pre-empting the need for onward movement of the *wh*-phrase. So not only does the presumed link between intermediate stop-overs and phase boundaries remain underived, so does the alleged need for movement to proceed via a succession of small steps.[35]

35 Of course 'EPP' could be divorced from particular probing features: it could be thought of as an autonomous property of a head. But as I pointed out earlier, that would weaken the bond between 'EPP' and probes or phase heads, and open the door to assignment of 'EPP' to functional heads that do not head a phase (or even to non-functional heads). In turn, that would give rise to a model of syntax in which the intermediate stop-overs along a movement path are no longer necessarily on the edges of phases: movement would then be a succession of steps, but the phase would play no part in it anymore.

The minimalist theory of syntax thus seems to have a hard time deriving in a principled manner the successive cyclicity of movement derivations that mainstream generative syntax has subscribed to for many decades. The bottom-up approach to the construction of syntactic structures and the movement operations taking place in them can deterministically ensure displacement to the edge of an intermediate in the course of long-distance movement only with an appeal to a custom-made ploy: the 'EPP property'. This property can be tied to phase heads only if 'EPP' is a property of a probing feature. If we do indeed link 'EPP' uniquely to the feature content of probes (phase heads), we must abandon the hope of deriving the successiveness of long movement from the labelling of [XP YP] structures, and we need to have recourse to feature inheritance to get the structural subject position occupied — i.e., to derive the effect of the original Extended Projection Principle. Though feature inheritance does not put 'derivation by phase' in peril (because the heads engaged in the inheritance relation are in the same phase), it does fundamentally weaken the bottom-up approach to syntactic structure building (see section 2.3.3).

What about the empirical support for successive-cyclic movement? Two things should be said about this. First of all, a detailed examination of the arguments for successive-cyclic movement available in the literature reveals that the vast majority of them are based on facts that are at best merely compatible with the hypothesis, not evidence for it (see Den Dikken 2009, to appear, and also the brief discussion of Van Urk & Richards 2015 at the end of chapter 4 of this work). Secondly, and more importantly in the context of a discussion of the directionality of syntactic structure building, to the extent that genuine evidence exists for the idea that long-distance filler–gap dependencies involve copies of the filler in intermediate positions along the way, such evidence is not an endorsement of the bottom-up approach. The contrary may very well be true: we have seen that the standard bottom-up model struggles to provide a decent rationale for successive-cyclic movement; the top-down alternative may actually fare better in motivating the existence of 'copies' in intermediate positions (see chapter 4). In section 2.5 I will examine one attempt to derive successive cyclicity from a top-down model (Chesi 2007).

2.3.5 Conclusion

Frequent assertions of the contrary notwithstanding, standard bottom-up minimalist derivations of syntactic structures do not proceed phase by phase. It is not the phase as a whole that gets handed over to the interpretive components when cyclic spell-out takes place: the edge of the phase must explicitly be exempted from being shunted to the interpretive components. If intermediate chunks of the structure get spelled out along the way, these are not in any obvious sense units that are useful to LF or PF (as Uriagereka 19xx and Epstein *et al*. 19yy observed early on). And if we need to keep track of the question of whether a particular feature, valued at stage S_n, was unvalued at an earlier stage S_{n-1} (recall the discussion of spell-out against the background of Chomsky 2001), then spelling S_{n-1} out does not allow us to 'forget' its component parts. So cyclic spell-out in a bottom-up derivation does not have any obvious advantages from the point of view of computational efficiency. In a top-down system, on the other hand, cyclic spell-out is not only feasible (with spell-out cycles defined in terms of maximal extended projections, *à la* (35)) but, as we have seen, also fully deterministic and efficacious.

Apart from the desirability of a top-down approach in connection with the cycle, one other important conclusion emerging from the discussion in this section is that domains for cyclic spell-out and syntactic opacity are not aligned. Different mechanisms are at play in delineating spell-out cycles and opaque domains — the former based on Grimshaw's (1991) notion of extended projection and the latter rooted in the Agree relation.

In the remainder of the discussion, I will have very little more to say about cyclic spell-out: I will assume it but will not contribute to it further. But opacity will play a leading role throughout the rest of the book. We will discuss in detail a wide variety of problems of locality in the realm of Ā-dependencies for which the top-down approach provides adequate solutions not available on alternative outlooks.

2.4 Working from the top down: Constituency conflicts and dependencies

At this point in the discussion, the main focus of this chapter (and of the book as a whole) shifts from the standard bottom-up approach to the alternative top-down model. In the remainder of this chapter, I will examine two extended pleas in the recent literature in favour of a top-down, left-to-right structure-building mechanism — the first, due to Colin Phillips, based on apparent constituency conflicts, and the second, advanced primarily by Cristiano Chesi, concentrating on the workings of filler–gap dependencies.

Before embarking on this exercise, let me forewarn the reader that reviewing these two extant arguments for a top-down, left-to-right grammar will take up quite a lot of space. The degree of detail in the discussion may perhaps be considered excessive. But it is important to dissect these arguments with care, because they are often either casually dismissed or blindly taken for granted in the literature. We need to ascertain whether the arguments truly show what they are supposed to show — and it will turn out, upon careful scrutiny, that they in fact do not.

2.4.1 Constituency conflicts: Phillips (1996, 2003)

Phillips (1996, 2003) was the first to champion a top-down left-to-right approach to structure building within a minimalist theory of syntax, explicitly with an eye towards bringing the grammar and the parser closer together. Applying a bottom-up generator of syntactic structure to on-line sentence processing is a major challenge (see esp. Stabler 2013 for discussion): sentences are not processed from the bottom right-hand corner (the last word or morpheme) to the top left-hand corner (the first word or morpheme) of the tree. So if syntactic structure building is to proceed from the bottom up, there need to be two separate models for the generation and processing of syntactic structures. While I agree that it would be desirable to have a single engine for both purposes, I will not be concerned here with questions regarding the efficacy of sentence processing. I will concentrate in my discussion of Phillips' syntax-internal arguments for top-down structure building.

Phillips' central concern is to account for discrepancies between various constituency tests and to explain why different structural diagnostics deliver different results. This concern has two specific subcomponents. I will discuss these in turn in the following subsections.

2.4.1.1 Temporary constituency

The first is that of the various constituency tests that we have at our disposal, some provide different results from others — in particular, coordination (Right Node Raising) often gives rise to very different assumptions about constituency than do movement and ellipsis. We see this, for instance, in the triples in (53) and (54). (On the poor status of pseudogapping with a direct object remnant in the double object construction, see Lasnik 1999:143.) Here coordination identifies something as a constituent that dramatically fails the VP-fronting and VP-ellipsis tests for constituency.

(53) a. John will [give his wife] and Bill will [give his mistress] dinner for Christmas
b. *[give his wife] though John will ___ dinner for Christmas, she won't be satisfied
c. *John will [give his wife] dinner for Christmas, and Bill will ___ lunch on her birthday

(54) a. John will [place the kettle on] and Mary will [take the kettle off] the stove
b. *[place the kettle on] though John will ___ the stove, Mary will not be satisfied
c. *John [place the kettle on] the stove, and Bill will ___ the kitchen counter

A significant portion of Phillips' work is devoted to showing that these results are not in fact in conflict with one another. The upshot of his discussion is that incremental top-down, left-to-right structure building can form temporary constituents at some point in the syntactic derivation which are destroyed in the subsequent structure-building process, and that coordination can take advantage of these temporary constituents where VP-fronting and VP-ellipsis cannot.

Thus, in the process of building the first conjuncts in (53a) and (54a), the syntax at some point delivers the strings *give his wife* and *place the kettle on* as constituents: just prior to the inclusion of the direct object in the first case and the object of P in the second, all we have in the VP is what is shown in (55a) and (55b) (where, in order to stay close to Phillips' work, I am assuming a 'Larsonian' approach to the structure of the VP).

(55) a. $[_{VP}$ *give* $[_{VP}$ *his wife* $[_{V'}$ <*give*>]]]
b. $[_{VP}$ *place* $[_{VP}$ *the kettle* $[_{V'}$ <*place*> $[_{PP}$ *on*]]]]

In the second conjuncts in (53a) and (54a), something gets added to these structures in the bottom right-hand corner: both *give* and *on* turn out to take a complement. Once this complement (*dinner*, *the stove*) is added at the bottom of the structure of the VP, we get (55′). Now we see that the strings *give his wife* and *place the kettle on* cease to be constituents.

(55′) a. $[_{VP}$ *give* $[_{VP}$ *his wife* $[_{V'}$ <*give*> [*dinner*]]]]
b. $[_{VP}$ *place* $[_{VP}$ *the kettle* $[_{V'}$ <*place*> $[_{PP}$ *on* [*the stove*]]]]]

For the coordination examples in (53a) and (54a) the fact that the constituency of *give his wife* and *place the kettle on* is destroyed is inconsequential: the syntax is welcome to add further material to the second conjuncts, to be shared with the first. (Phillips defends a multi-dominance approach to Right Node Raising, which we need not go into here.) But for the VP-fronting and VP-ellipsis examples in (53b,c) and (54b,c), we run into trouble. What makes the b– and c–examples different is that they both contain a gap (marked with '___') that the antecedent-VP needs to identify and license. In the following paragraphs, I will briefly summarise the account of why the b– and c–examples are ungrammatical.

Assume, as is entirely standard, that a non-pronominal empty category needs to be licensed by a constituent serving as its antecedent. This allows us to understand what is going wrong in the c–examples in (53) and (54), involving ellipsis: although the strings *give his wife* and *place the kettle on* are temporarily identified as constituents in the course of the composition of the first conjuncts of these examples, their constituency vanishes once *dinner* and *the stove* are integrated into the structure of those initial conjuncts. By the time we are confronted with the ellipsis site and need to identify a constituent as its antecedent, we no longer have *give his wife* and *place the kettle on* as constituents. So the c–examples fail because of the fact that the bracketed strings in these sentences are no longer constituents at the point at which the ellipsis site is to be licensed.

For the VP-fronting examples in (53b) and (54b), the constituency of the fronted constituents *give his wife* and *place the kettle on* is preserved throughout the derivation. So why can't *dinner* and *the stove* be added at the bottom right of the structure of the reconstructed VP? Phillips (2003:§4.6.2) wants to derive the ungrammaticality of examples of the type in (53b) and (54b), and also those in (53c) and (54c), from what a constraint that says that VP-ellipsis and VP-fronting may target only 'potential complete VPs'. Since *give his wife* and *place the kettle on* are not 'potential complete VPs', the outputs in (53b,c) and (54b,c) are ill-formed because they contain subportions (viz., the bracketed strings) that fail the constraint. For Right Node Raising, this constraint is not active because Phillips assumes that in RNR constructions we are dealing with a multi-dominance structure, with the element following the second bracketed string in the a–examples being shared by both conjuncts; so the VP of the first conjunct is not incomplete.

The 'potential complete VP' constraint on VP-fronting and VP-ellipsis is designed to block outputs of these operations in which an obligatory argument is stranded. But minor changes in the examples in (53) and (54) show that the facts remain exactly the same when the fronted or elided strings ARE in fact 'potential complete VPs' (see also Landau 2007). The bracketed strings in the examples in (53′) and (54′) are well-formed as such: they serve as predicates of the grammatical sentences *John will serve his wife*[36] and *John put the kettle on*.[37] The fact that the bracketed strings in (53′) and (54′) are complete and well-formed VPs should make them legitimate antecedents for the gaps in the b–examples in (53) and (54). Nonetheless, the VP-fronting and VP-ellipsis examples are as bad as their counterparts in primeless sentences above.

(53′) a. John will [serve his wife] and Bill will [serve his mistress] dinner for Christmas
b. *[serve his wife] though John will ___ dinner for Christmas, she won't be satisfied
c. *John will [serve his wife] dinner for Christmas, and Bill will ___ lunch on her birthday

(54′) a. John will [put the kettle on] and Mary will [take the kettle off] the stove
b. *[put the kettle on] though John will ___ the stove, Mary will not be satisfied
c. *John will [put the kettle on] the stove, and Bill will ___ the kitchen counter

As far as I can see, the fact that the VP-fronting examples in (53b′) and (54b′) are ungrammatical will follow only if we require strict parallelism between '___' and its antecedent regarding their internal syntactic structure. Once the direct object in (53b,b′) and the P-object in (54b,b′) are integrated into the structure of the complement of *will*, the constituency of the strings *give/serve his wife* and *place/put the kettle on*, recovered in the position of '___', dissolves. As a result, the necessary parallelism between the antecedent, for which the constituency of these strings IS preserved, and the gap fails to materialise: the internal structures of the fronted VP and the gap are different, with the former lacking an object altogether and the latter featuring an object 'tucked in' at the bottom right, below the lowest head.

36 Those readers who find *John will serve his wife* incomplete should feel free to replace the text examples involving *serve* with variants in which the subject is a noun phrase headed by *waiter* (where unspecified object deletion in the case of *serve* is particularly felicitous: *this waiter serves table 1*) — for instance, as follows: *this waiter will serve table 1 and that waiter will serve table 2 dinner this evening*.

37 The word *on* in *put the kettle on* is arguably a transitive preposition with an implicit object (rather than an intransitive particle).. Even if this implicit object is represented in syntax (as a silent noun phrase), the b–example in (54′) should be grammatical, with the implicit object as a hyperonym and *the stove* as a hyponym (cf. *as for fish, I like flounder*).

Upon close scrutiny, therefore, we find that the temporary constituents delivered by a top-down, left-to-right structure-building model of the type proposed by Phillips (1996, 2003) do not turn out to sustain 'destruction' under further structure building when such additional structure building involves the addition of an argument: even non-obligatory arguments strongly resist stranding under pseudogapping in double object constructions and under VP-fronting. The only temporary constituents that survive 'destruction' under the addition of an argument are those found in Right Node Raising constructions. But in the multi-dominance analysis advocated by Phillips, the RNR'ed phrase is in fact a constituent of the first conjunct. So it is not clear that we are dealing with a genuine 'temporary constituent' in the coordination cases. At any rate, the analysis of RNR is sufficiently uncertain to make any claims about the efficacy of a particular structure-building model based on RNR inherently weak. The impression emerges that there are no temporary constituents produced by a top-down, left-to-right structure-building mechanism that could possibly be destroyed by the addition of an argument later in the derivation.

That said, we are left to investigate whether temporary constituents could perhaps be altered by the addition, at the bottom right of the structure, of *non*-arguments. It is at this point that we need to turn our attention to the second major ingredient of Phillips' account of conflicts in constituency. This is the topic of the next subsection.

2.4.1.2 C-command and binding: 'Pesetsky paradoxes'

The second component of Phillips' (1996, 2003) discussion of constituency conflicts is what he calls 'Pesetsky's paradox' (after Pesetsky 1995): the fact that a constituent that does not, on orthodox assumptions, c-command an element that, again on orthodox assumptions, demands a c-commanding antecedent can nonetheless serve as an antecedent for that element. It is here that we find some striking discrepancies between VP-fronting and VP-ellipsis.

To set up the key cases, let us first take a look at the baseline examples that led Pesetsky (1995) to set up what he called 'cascade' structures for certain verb phrases. In (56) and (57), we see an apparent conflict between what a standard 'layered' syntactic structure would lead us to expect for the organisation of the VP and what the dependencies between *each other* and *his* and their antecedents (*them* and *everyone*) would demand. On standard assumptions, the PPs *in the garden* and *on each other's birthdays* in (56), and *in a speech* and *at his graduation* in (57), are right-adjoined to VP or higher, while *to them* and *everyone* are contained in V′. Equally standard, however, is the assumption that an anaphor (here the reciprocal *each other*) or a bound-variable pronoun (*his*) needs to be bound by a c-commanding antecedent. Unless one of these standard assumptions is wrong, this creates a paradox.

(56) John gave books to *them* in the garden on *each other*'s birthdays
(57) the principal congratulated *everyone* in a speech at *his* graduation

Pesetsky argues that while the second assumption is correct, the first is not: although layered, left-branching structures do exist, it is possible for the modifier-PPs in the examples in (56) and (57) to be merged low in the structure, within the c-command domain of *them* and *everyone*. When they are merged low, we arrive at strikingly unorthodox structures in which every incoming phrase is integrated into the preceding structure in a right-branching structure. Thus, for (56) we get (58b) as its cascade structure, and for (57) we arrive at (59b). The traditional left-branching layered structures (infused, for (56), with a Larsonian approach to ditransitives) are given for these sentences in (58a) and (59a), respectively.

(58) a. [$_{VP}$ [$_{VP}$ [$_{VP}$ *give* [$_{VP}$ *books* [<*give*> [$_{PP}$ *to them*]]]] [$_{PP}$ *in the garden*]] [$_{PP}$ *on e.o.'s birthdays*]]
b. [$_{VP}$ *give* [$_{PP}$ *books* [$_{P'}$ *to* [$_{PP}$ *them* [$_{P'}$ *in* [$_{PP}$ *the garden* [$_{P'}$ *on* [*e.o.'s birthdays*]]]]]]]]

(59) a. [$_{VP}$ [$_{VP}$ [$_{VP}$ *congratulate* [*everyone*]] [$_{PP}$ *in a speech*]] [$_{PP}$ *at his graduation*]]
b. [$_{VP}$ *congratulate* [$_{PP}$ *everyone* [$_{P'}$ *in* [$_{PP}$ *a speech* [$_{P'}$ *at* [*his graduation*]]]]]]

The cascade structures are very useful for accounting for the binding facts in (56) and (57). These are now perfectly straightforward: *them* c-commands *each other*, and *everyone* c-commands *his*. But cascade structures are not particularly helpful elsewhere in the grammar. They would seem to make it difficult to account for the fact that the inner portions of the VP can be fronted or elided separately, stranding the adverbial modifiers, as in (60) and (61). For these kinds of examples, a traditional layered structure for the VP would come in handy. Pesetsky argues, therefore, that cascade structures are not the only ones that are assignable to complex VPs of the type found in (56) and (57): these co-exist with layered structures of the familiar sort.

(60) a. John said that he would give books to them, and [give books to them] he did ___ in the garden at Christmas
b. John gave books to them in the garden at Christmas, and Mary did ___ in her office on the first day of the school year

(61) a. the principal said that he would congratulate them, and [congratulate them] he did ___ in a speech at graduation
b. the principal congratulated them in a speech at graduation, and the secretary did ___ in a letter on the first day of the school year

The two structures must exist side by side, because otherwise the facts of VP-fronting once combined in a single sentence with the binding facts presented in (56) and (57) seem to become irreconcilable. The examples in (62) and (63) are grammatical (regardless of whether *in the garden* and *in a speech* are included in the fronted VP or not).

(62) John said that he'd give books to *them* (in the garden), and
a. [give books to *them* in the garden] he did ___ on *each other*'s birthdays
b. [give books to *them*] he did ___ in the garden on *each other*'s birthdays

(63) the principal said he'd congratulate *everyone* (in a speech), and
a. [congratulate *everyone* in a speech] he did ___ at *his* graduation
b. [congratulate *everyone*] he did ___ in a speech at *his* graduation

Pesetsky's cascade structure for the VP, which is totally right-branching, cannot serve as the input to VP-fronting as in (62) or (63) — the bracketed strings in these sentences are not constituents in (58) and (59). Applying VP-fronting to a layered structure more familiar from traditional X-bar Theory would solve the constituency problem but would not make it possible for *them* to c-command *each other* or for *everyone* to c-command *his*. If binding requires c-command, and c-command in these examples requires a cascade structure, but only a layered structure can deliver the output seen in the partial VP-fronting examples in (62) and (63), then we find ourselves confronted with a paradox. Pesetsky resolves it by assuming that for every VP there are two parallel structures, one layered and the other a cascade, and that syntactic operations such as VP-fronting target the layered structure while binding applies to the cascade structure.

Note that it is not just for reasons having to do with constituency that we cannot content ourselves with just the cascades structure in (58b) and (59b). These structures also do not serve the purposes of θ-role assignment very well. In (58b) we have trouble associating *them* to the Goal θ-role assigned to it (depending on one's assumptions) either by the verb *give* or by the preposition *to*, and we also do not manage to get *in* to assign a θ-role to *the garden*. And in (59b), with *everyone* occupying the specifier position of the PP projected by *in* and *a speech* sitting in the specifier of the PP projected by *at*, we cannot link these arguments to their θ-role assigners under sisterhood. For Pesetsky, therefore, cascades could never be the sole syntactic structures assigned to sentences of the type in (56) and (57): a layered structure must be available for the purpose of argument structure representation.

Phillips (1996, 2003) finds an ingenious way around the θ-role assignment problem. He marries the Pesetskian cascade to the Larsonian VP-shell, and comes up with the representations in (58b′) and (59b′) as the right-branching cascade-type structures for the simplified examples in (56′) and (57′).[38]

(56′) John gave books to *them* on *each other*'s birthdays
(58b′) [$_{V'}$ *give* [$_{VP}$ *books* [$_{V'}$ <*give*> [$_{PP}$ *to* [$_{VP}$ *them* [$_{V'}$ <*give*> [$_{PP}$ *on* [*e.o.'s b'days*]]]]]]]]

(57′) the principal congratulated *everyone* at *his* graduation
(59b′) [$_{V'}$ *congratulate* [$_{VP}$ *everyone* [$_{V'}$ <*congratulate*> [$_{PP}$ *at* [*his graduation*]]]]]

Now *books* and *them* are represented as arguments of *give* (which is represented in the structure three times, but spelled out only in the highest V-slot), each generated in VP-specifier positions; and *everyone* is likewise represented as an argument of *congratulate*, originating in the specifier of the lower VP. The PPs containing the bound elements are in the complement position of the lowest V-head in the structure, and still c-commanded, as desired, by their binders.

Before proceeding, let me draw attention to a non-trivial detail that potentially undermines the efficacy of Phillips' Larsonian revisions of Pesetskian cascade structures. Note that in (58b′) the PP projected by *to* interrupts the VP-shell structure of *give*, which is resumed below *to*. If, as Larson (1988) argues, the various V-heads in a VP-shell structure are linked to one another in a chain (i.e., via movement), *to* in (58b′) interrupts the chain that is supposed to link up the three copies of *give*. This matter is far from trivial. It will need to be satisfactorily resolved before we can conclude with confidence that the possibility of representing Pesetsky-type sentences exclusively in terms of a right-branching cascade structure is real.

38 The structure in (58b′) is a hybrid of Pesetsky's cascades and Larsonian VP-shells. The 'pure' Pesetskian cascade is in (58b); the standard Larsonian VP-shell structure for the string in (56) would look as in (i), which does not serve Phillips' purposes because it does not allow the PP-contained antecedent of the reciprocal to c-command the latter: the PP-node projected by *to* that dominates the binder and excludes the bindee precludes c-command.

(i) [$_{V'}$ *give* [$_{VP}$ *books* [$_{V'}$ <*give*> [$_{VP}$ [$_{PP}$ *to* [*them*]] [$_{V'}$ <*give*> [$_{PP}$ *on* [*e.o.'s b'days*]]]]]]]

Note that I kept the reworkings of cascade structures in (58b′) and (59b′) simple by looking only at versions of (56) and (57) containing just one adverbial PP. But Phillips' proposal can account for the more complex cases as well, maintaining local θ-role assignment throughout: thus, for the portion of the structure spelled out as *in the garden on each other's birthdays*, a 'Larsonian PP' can be projected, with *the garden* in the specifier of the lower *in*-PP, locally θ-marked by *in*:

(ii) [... [$_{V'}$ <*give*> [$_{PP}$ *in* [$_{PP}$ *the garden* [$_{P'}$ <*in*> [$_{PP}$ *on* [*e.o.'s birthdays*]]]]]]]

But if it can be done, this will certainly be progress, for it would seem rather awkward, from a conceptual point of view, for the grammar to assign two complete structures to every VP. Moreover, Phillips (1996, 2003) argues that from an empirical perspective the parallel structures approach is also not the right way to proceed. He notes that while VP-fronting allows a PP containing an anaphor or bound-variable pronoun to be stranded in the clause, VP-ellipsis does not allow such a PP to be orphaned outside the ellipsis site. We see this in (64) and (65).

(64) *John gave books to *them* in the garden on *each other*'s birthdays, and Mary did ___ (in her office) on *each other*'s first day of school

(65) *the principal congratulated *everyone* in a speech at *his* graduation, and the secretary did ___ (in a letter) on *his* birthday

So when VP-ellipsis targets a subpart of the complex verb phrase and leaves the PP containing the binding-dependent element behind, binding actually fails.

Phillips' (1996, 2003) major contribution to the puzzle posed by the facts reviewed in the preceding paragraphs is the argument that top-down left-to-right structure building makes the contrast presented by the VP-fronting cases and the VP-ellipsis cases fall out. Let us examine this argument closely to see if it holds up to careful scrutiny.

I will start with VP-fronting. On a top-down approach, the structure of VP in the left periphery is built first (probably in a right-branching way: Phillips assumes that this is the preferred option; but it does not actually matter what the internal structure of the fronted VPs in (53) and (54) looks like), and later in the derivation the structure of the matrix VP (in the position of the gap, '___') is reconstructed based on the material in the antecedent (i.e., the fronted VP). Crucially, additional incoming material belonging to the VP can be accommodated below the last reconstructed element, 'added to the right of the reconstructed VP, at the bottom of the right-branching VP' (Phillips 2003:52). So for the examples in (62) and (63), the top-down approach allows the adverbial modifiers *on each other's birthdays* and *at his graduation* to be merged low to the reconstructed copy of the fronted VP in the position of the gap, at the bottom-right edge of the structure. Merged this low, and with the reconstructed VP having a cascade structure, the bound elements (*each other*, *his*) end up within the c-command domain of their binders (*them*/*everyone*). This is something that we would not be able to achieve in a bottom-up derivation, on a movement account of VP-fronting:[39] if we started out with a cascade-structured VP in the base position, we would never be able to isolate a movable constituent that would exclude the PPs containing the bound expressions. We would be able to front the entire VP including these PPs, but we would not be able to strand them.

That VP-ellipsis is different from VP-fronting with regard to stranding of material containing an element to be bound by something in the elliptical VP also falls out from the top-down, left-to-right approach. In (64) and (65), if *on each other's birthdays* and *at his graduation*

39 A realistic alternative, however, would treat VP-fronting in terms of base-generation. Such an approach is potentially supported for English by the fact that the fronted VP can fail to show morphological connectivity with the selecting verb: when the fronted VP is the semantic dependent of the auxiliary of the perfect, *have*, which ordinarily selects a past-participial constituent, we find that the verbal head of the fronted VP can be a bare infinitive rather than a past participle (Emonds 1976):

(i) if he hadn't died, they would have had to reassign him because [resign], he never would have *ec*

A base-generation analysis of VP-fronting can accommodate the bare infinitive in brackets in (i) as the default or citation form of the verb. A movement analysis would have a much harder time accounting for the form of the verb.

in the first conjunct were to be merged low, in a right-branching cascade structure, these PPs would be reconstructed into the ellipsis site, yielding #*Mary gave books to them on each other's birthdays on each other's first day of school* and *the secretary congratulated everyone at his graduation on his birthday*, which are not the intended interpretations (if they make sense at all). So to avoid such an unwanted outcome of reconstruction of the elliptical VP, the antecedent-VP must have a layered structure, with *on each other's birthday* and *at his graduation* attached as right-peripheral adjuncts to VP or higher. In such a left-branching structure of the traditional type, *them*/*everyone* cannot c-command *each other*/*his*, so the binding-theoretic requirements imposed on the reciprocal/bound-variable pronoun are not met. This predicts correctly that while the first conjuncts of (64) and (65) in and of themselves are perfectly grammatical, the coordination-*cum*-ellipsis constructions as a whole are not. And it also explains why VP-ellipsis is different from VP-fronting when it comes to the relevant binding facts.

These good results emerging from Phillips' (1996, 2003) account of 'Pesetsky's paradox' notwithstanding, however, his argument for top-down left-to-right structure building based on the contrast between (62)/(63) and (64)/(65) is nonetheless empirically flawed. The problem lies primarily in the account of the VP-ellipsis facts. The examples in (66) and (67) (below) are key in this. Let us consider first the example in (66), which is perfectly fine (as Phillips 1996 acknowledges, though his 2003 paper is silent on this; see also Lechner 2003), in contrast to (64).

(66) John gave books to *them* on *each other*'s birthdays, and Mary did ___ at Christmas

By the logic of the discussion of (64), the presence in the structure of a temporal PP outside the ellipsis site that is contrasted with a temporal PP in the first conjunct should force the latter PP into a left-branching structure: otherwise we would necessarily reconstruct the temporal PP in the first conjunct into the elliptical VP in the second conjunct, yielding the unwanted #*Mary gave books to them at Christmas on each other's birthdays*. With *on each other's birthdays* in (66) thus necessarily outside the node containing *give books to them*, we cannot establish the desired c-command relationship between *them* and *each other*. The sentence in (66) is thus expected to be ill-formed, on a par with (64), contrary to fact. The grammaticality of (66) creates a problem for the proposal.

Phillips & Lewis (2013:39) make a 'speculative' attempt to save the account for (66) by appealing to 'incremental interpretation': 'Once the binding relation in the first clause of [(66)] is established, it cannot be retracted based on the constituency requirements of the VP-ellipsis construction'. It seems that what they have in mind here is that we can allow the grammar/parser to construct a cascade structure for the first conjunct, with binding of *each other* being taken care of once and for all, and that even when this structure must subsequently be revised once the ellipsis is resolved ('based on the constituency requirements of the VP-ellipsis construction'), the binding relation established in the rejected right-branching parse is preserved.

If we assume this speculation of Phillips & Lewis (2013), it seems to lose us one of the advances of Phillips' (1996, 2003) reworking of Pesetskian cascade-*cum*-layered syntax — the desirable abandonment of parallel structures is put in jeopardy by the revision of the right-branching structure of the VP in the first conjunct. Since Phillips assumes the grammar and the parser to form a single engine, this also requires the parser to either be non-deterministic or consider multiple possible parses in parallel ('parallel parsing'). But most seriously, from a syntax-internal point of view, the speculation about incremental interpretation makes Phillips' (1996, 2003) analysis of examples of the type in (64) (repeated below) null and void. Let us see why.

(64) *John gave books to *them* on *each other*'s birthdays, and Mary did ___ on *each other*'s first day of school

Recall that Phillips rules out (64) by saying that *each other* in the first conjunct lacks a c-commanding antecedent because the *on*-PP that it is in must necessarily be merged as a right-hand adjunct to VP in order to be prevented from being copied into the ellipsis site. If we now revise the theory so as to temporarily allow a right-branching structure for the VP in the first conjunct for binding purposes, we lose the account we had for (64). The right-branching cascade structure for the VP in the first conjunct will give *each other* its c-commanding antecedent, with binding sanctioned — once and for all. When we eventually discover that *on each other's birthdays* needs to be a VP-level adjunct in order for the ellipsis site to be properly identified, this should have no consequences for binding of *each other* in the first conjunct. Since *on each other's first day of school* in the second conjunct can be merged in at the bottom of a right-branching VP, that anaphor, too, can be bound by a c-commanding antecedent: the reconstructed VP *give books to them* can be perfectly right-branching, just like its identifier. The suggestion that once binding in the first conjunct is established, it cannot be withdrawn thus seems to undermine the successful account of (64). It seems that we cannot get (64) ruled out while at the same time getting (66) ruled in.

Alongside (64), consider its inverse in (67), with the anaphor-containing temporal PP outside the ellipsis site in the second conjunct, and *at Christmas* as the temporal PP in the first.

(67) *John gave books to *them* at Christmas, and Mary did ___ on *each other*'s birthdays

For (67) (minimally adapted from Phillips & Lewis 2013), we need *at Christmas* to serve as an adjunct to VP, in order for it not to be included in the copy reconstructed into the ellipsis site. The right-adjunct status of *at Christmas* has no adverse consequences at all for the first conjunct: there is no element inside this PP that is in need a c-commanding antecedent. So the first conjunct of (67) should be perfectly fine with *at Christmas* right-adjoined to the VP. The rest of the VP is free to be either left-branching (with *books* in the complement of V and *to them* adjoined to VP) or right-branching (with *to them* as the complement of V in a Larsonian shell structure, and *books* as the specifier of that VP). For the structure of the first conjunct, the choice between these options is immaterial. A VP structure in which one part (*give books to them*) is right-branching and another part (the one created by right-adjunction of *at Christmas*) is left-branching is perfectly coherent. And with the copy of *give books to them* in the ellipsis site assigned a right-branching Larsonian structure, we should be able to derive a grammatical result in the second conjunct, once *on each other's birthdays* is added at the bottom of the right-branching VP: *each other* should be able to be bound by *them* under c-command in the cascade structure of the second conjunct. We thus expect (67) to be grammatical. But Phillips & Lewis (2013) themselves observe that binding actually fails in this environment.

(66) and (67) combined seem to suggest that Phillips has misdiagnosed the problem with (64). When we put these three examples together in one block, we quickly discover what the root of the binding problem is, descriptively speaking:

(64) *John gave books to *them* on *each other*'s birthdays, and Mary did ___ on *each other*'s first day of school
(66) John gave books to *them* on *each other*'s birthdays, and Mary did ___ at Christmas
(67) *John gave books to *them* at Christmas, and Mary did ___ on *each other*'s birthdays

The generalisation that emerges is that in a VP-ellipsis construction the reciprocal can be bound legitimately only if it has an overt antecedent within the same clause. (The same holds, *mutatis mutandis*, for bound-variable anaphora.) We know that there is no blanket ban on 'orphaned' bindees: in VP-fronting constructions, the result of stranding a PP containing a bindee is grammatical, as (62) showed for *each other*.

(62) John said that he'd give books to *them* (in the garden), and
 a. [give books to *them* in the garden] he did ___ on *each other*'s birthdays
 b. [give books to *them*] he did ___ in the garden on *each other*'s birthdays

So why are VP-fronting and VP-ellipsis different with respect to this 'orphaning'? What I would like to suggest here is that we can explain the difference between the two processes in this realm directly on the basis of their mechanics, independently of the question of whether syntactic structures are built from the top down or from the bottom up — and that therefore, the stranding facts fail to adjudicate between the two competing approaches to the directionality of syntactic structure building.

Observationally, the key difference between VP-fronting and VP-ellipsis is that the former is (or at least, can be; recall fn. 39) a movement operation, leaving a copy behind that is identical with its antecedent, while the latter is not. In both construction types, there is silence in the position marked by '___'; but only in a movement derivation of VP-fronting is that silence necessarily fully identical with (i.e., a literal copy of) the overt VP. For VP-ellipsis, what we know from the literature is that there needs to be semantic parallelism between the ellipsis site and the antecedent. But whether the elliptical VP includes all of the ingredients of the antecedent remains an open question. Let us investigate this question with reference to the cases under discussion.

In VP-ellipsis constructions, the identifier of a VP-ellipsis site (the antecedent-VP) may very well contain information that is new in the discourse. But the elliptical VP whose content is recovered by the antecedent-VP can contain only old information, anaphoric to the information expressed by the antecedent-VP. We know that natural language prefers not to express old information if it can get away with suppressing it: this is precisely why ellipsis is such much more natural than pronunciation of the recoverable VP in examples like the ones we have been discussing. Let us revisit the example in (60b), repeated below as (68a), for illustration. If instead of eliding the VP in the second conjunct we had pronounced all of the content of the matching VP in the first conjunct, the result (given in (68b)) would have been pragmatically rather odd (though of course by no means ungrammatical).

(68) a. John gave books to them in the garden at Christmas, and Mary did ___ in her office on the first day of the school year
 b. #John gave books to them in the garden at Christmas, and Mary gave books to them in her office on the first day of the school year

But not only do we expect pragmatic preferences to make (68b) unnatural compared to (58a), we also expect that recovery in the ellipsis site of any and all material that is informationally redundant and syntactically and semantically dispensable would be dispreferred. Of course the verb will have to be recovered in the ellipsis site: that is a structural requirement without whose fulfilment the subject of the second conjunct could not be licensed. But besides the verb (the predicate head), nothing is strictly required to be recovered in the ellipsis site in the

examples under investigation. In (68a), along with the recovery of the verb *give*, we get a θ-role for *Mary*, and we deduce, from our knowledge about 'giving' events, that there will be a thing given and a beneficiary involved in the event as well. The thing given and the beneficiary are both identified already in the antecedent-VP, as *books* and *them*: they are known. Semantically, there is no reason, therefore, to assume that anything besides *give* is reconstructed into the ellipsis site. And syntactically, we know that *give* can survive on its own as well, as, for instance, in the conversation in (69), and in (70) (the latter an attested example culled from the internet).

(69) A: would you care to give some money to those in need this holiday season?
B: I gave already

(70) the usual response when I knocked on someone's door to ask for money was 'I gave already'

And of course, in both examples the interlocutor could also have used the lexical verb *do* in combination with a pronominal object (a 'deep anaphor' in the sense of Hankamar & Sag 197x), as in (71):

(71) I did it/that already

The smallest piece of verbal structure that we can get away with in the ellipsis site in the examples at hand, therefore, is a projection of *give* or *do* with an external argument and an implicit or pronominal complement. (I am assuming that in (69B) and (70), *give* has no syntactically projected implicit arguments; the discussion of Unspecified Object Deletion dates back at least to Chomsky 1965:87.) Assuming that less is best (i.e., that we do not postulate anything 'old' in the ellipsis site beyond the bare minimum), we arrive at the conclusion that in the VP-ellipsis cases under discussion earlier in this section, all that we have in the ellipsis site is *give* or *do it*. Let us reproduce the examples in (64), (66) and (67) with the ellipsis site recovered as such:

(64′) *John gave books to *them* on *each other*'s birthdays, and Mary did [$_{VP}$ give/do it] on *each other*'s first day of school

(66′) John gave books to *them* on *each other*'s birthdays, and Mary did [$_{VP}$ give/do it] at Christmas

(67′) *John gave books to *them* at Christmas, and Mary did [$_{VP}$ give/do it] on *each other*'s birthdays

The ungrammaticality of (64′) and (67′) is easy to understand: in the second conjunct, there is no linguistically represented antecedent for *each other* in the second conjunct at all. That (66′) is fine is also straightforward: the bare VP in the second conjunct is not inconveniencing *at Christmas* in any way; it can happily modify this VP (*I already gave/did it at Christmas* is perfectly fine in contexts of the type in (69) and (70)).

So if we assume that VP-ellipsis reconstructs as little as possible into the ellipsis site, we can easily understand why the examples in (64) and (67) are ungrammatical while (66) is well-formed. We can do this entirely without any appeals to the direction of structure building: all we need is the recognition of the fact, known independently, that repetition of old information is dispreferred in natural language. For ellipsis, which is itself a strategy used in response to the desire to suppress redundancy, this dispreference naturally results in the complete omission, in the

process of reconstruction of the ellipsis site, of any and all material that is not required for syntactic or semantic reasons.[40]

Let us at this point make an interim summary of what we have seen. For Phillips, the contrast between VP-ellipsis and VP-fronting with regard to the stranding of constituents harbouring material bound by something in the elliptical or fronted VP presented an argument for a top-down, left-to-right approach to the way syntactic structures are built. Such an approach delivers the difference between (62)/(63), on the one hand, and (64)/(65), on the other, by forcing *on each other's birthdays* and *at his graduation* in the latter examples into a right-adjoined position, outside the c-command domain of the binders of the anaphor and bound-variable pronoun. A right-branching cascade structure is unavailable for the antecedent-VP in these examples because it would make the reconstructed VP in the gap position semantically too rich. But by the logic of right-branching bias of Phillips' top-down approach, the stranded material in the second conjunct, to the right of the gap, should be welcome to be added in the bottom right-hand corner of the reconstructed VP. This leads to the expectation that (67) ought to be perfectly well-formed, which it is not. And if the counterpart of the stranded PP in the second conjunct should always be attached as a right-adjunct in the *first* conjunct, (66) should be ungrammatical, but it is fine. So Phillips' account of (64)/(65) fails to make the right predictions. There is a much simpler analysis available — and that simpler account (which assumes that the reconstructed VP contains just as much material as is necessary for convergence) does not depend on the direction of structure building or, for that matter, on the question of whether syntactic structures are right-branching cascades or left-branching layered structures. There is, it turns out, no 'Pesetsky paradox' in VP-ellipsis constructions at all.

Now let us now return to the VP-fronting constructions in (62) and (63). For these, Pesetsky and Phillips have likely overestimated the severity of the problem. One should point out (as do Pollard & Sag 1992, Janke & Neeleman 2012, and Bruening 2014) that when serving as the possessor of a noun phrase, the reciprocal *each other* does not show the typical behaviour of an anaphor: it behaves more like a logophor, allowing a non-local antecedent:

(71) a. **they* hoped that the psychologist would tell *each other* their secrets
b. *they* hoped that the psychologist would tell *each other*'s secrets to them

One should also mention the fact (noted by Kayne 1994 and many others) that *everyone* seems rather good at binding a variable pronoun in syntactic environments in which traditional definition of c-command would not readily deliver a c-command relation between binder and bindee. Thus, in (72), *every little boy* can bind *his*, which is outside the complex noun phrase that, on standard assumptions, dominates its possessor.

(72) [*every little boy*'s mother] is *his* favourite future wife

40 Though I couched the above discussion of VP-ellipsis in 'reconstruction' terms (following Phillips' practice), this account of the facts reviewed does not actually depend on any particular choice of approach to ellipsis: in particular, it is compatible both with an analysis that assumes an unstructured proform that gets recovered at LF and with an approach that says that the elliptical constituent has internal structure throughout the syntactic derivation and is silent because it does not get a phonological matrix at PF. Both analyses are compatible with the proposal because regardless of which of the two overall perspectives we take, it will always be possible to postulate just [$_{VP}$ *give*] as the elliptical VP — as a recovered/reconstructed minimal VP at LF, or as an underlying minimal VP throughout the syntax. All that matters for the purposes of explaining the facts is that in VP-ellipsis constructions (and presumably in ellipsis constructions across the board) the content of the ellipsis site is not allowed to be richer than the bare minimum needed for grammatical convergence.

Assuming, as is standard, that universal quantifiers must undergo movement out of their θ-position ('Quantifier Raising'), we may naturally hypothesise that a universal quantifier extends its c-command domain all the way up to the node to which QR attaches it. For quantifiers originating in object positions, that node is minimally as large as *v*P. It should not be difficult, therefore, for a universal quantifier to antecede a bound-variable pronoun in a PP adjoined to VP in a traditional left-branching layered structure.[41]

41 Bruening (2014) argues that much more generally, the c-command domain for elements contained in the VP is as large as *v*P, the first phase dominating the VP. He argues for a definition of command that makes key reference to the phase:

(i) *phase command* (Bruening 2014)
X phase-commands Y iff there is no ZP, ZP a phasal node, such that ZP dominates X but does not dominate Y

If, as Bruening has it, PP is not a phase, then the first phase that dominates a PP-contained element embedded inside the VP will be the *v*P, which will allow the complement of P to command other material dominated by *v*P. If correct, this will take care of the 'Pesetsky paradox' even if *each other* in the possessor position of a DP should need a local binder that commands it (*pace* (71b)); and of course it accommodates the variable pronoun binding facts as well, without reference to QR being necessary. With phase command in place, and with binding defined in terms of it, the need for cascade structures evaporates, and standard left-branching layered structures will cover not just the preposing and ellipsis facts but also the binding data.

We should carefully assess, however, whether it is useful and desirable to stretch the command domain up all the way to the phase. The usefulness of this move will be severely diminished if it should turn out that PPs can have phasal status after all. This is rejected explicitly by Bruening, but his discussion does not make mention of the detailed investigations of the functional extended projection of P in the recent literature (Svenonius 2010, Den Dikken 2010, Aelbrecht & Den Dikken 2013). And the desirability of stretching the command domain up to the phase is jeopardised by the need to make reference to precedence in a definition of binding based on phase command:

(ii) A binds B iff A and B are coindexed and A precedes and phase-commands B

This return to precede-and-command conjures up the same kinds of facts discussed by Reinhart (1976) in defence of strict c-command and against reference to precedence. Bruening manages to handle a large number of these facts with an account using reconstruction and late merger. But a small but significant core of data does not lend itself to such an approach. Particularly recalcitrant are sets such as (iii)–(iv) (from Reinhart 1976:23ff.):

(iii) a. **he* is very obnoxious to *Ben*'s friends
b. *I found *him* to be very obnoxious to *Ben*'s friends
(iv) a. *how obnoxious to *Ben*'s friends *he* is!
b. how obnoxious to *Ben*'s friends I found *him* to be!

The peculiar thing about (iv) is that we get a Principle C effect with *he* qua matrix subject in (iva) but not with *him* qua ECM-subject of the infinitival clause in (ivb). Reinhart has an account for this based on her definition of c-command, crucially eschewing reference to precedence: *he* c-commands, hence inadvertently binds, *Ben* in (iva) but not in (ivb). No alternative account of these data is available to my knowledge. A definition of binding in terms of phase-command *cum* precedence, along the lines of (ii) certainly will not manage to accommodate (iv). (The facts in (iv) are also highly problematic for Huang's 1993 proposal regarding the lack of binding ambiguity under fronting of predicates containing anaphors.)

See also Zwart (2015) for a critical discussion of Bruening's (2014) 'phase command' in (I). It would lead me too far afield to delve into the (de)merits of phase-based approaches to command at more length here. But the reader can easily verify that on the approach to phases developed in this monograph, phase command would not be a workable alternative to c-command, leading to vastly overstretched a command domain in many cases.

Bearing these two remarks in mind, we are immediately absolved of the need to say anything special about the VP-fronting examples in (62) and (63), or, for that matter, about the baseline cases in (66) and (67): traditional layered structures will do the job for all of these. Of course, if we set the *each other*-as-possessor facts aside with an appeal to (71), it will still be incumbent on us to explain the ungrammaticality of (64). (Logophoricity will not help us out here: the antecedent of the reciprocal is quite clearly not the logophoric centre of these sentences.) And if we respond to (72) by allowing quantifiers a wider c-command range in general, so that cascade structures might become superfluous for the account of (67), we will still need to explain the fact that binding breaks down in (65). But recall that we actually have a simple account of why (64) and (65) are ill-formed — one that, unlike Phillips' proposal, does not make the wrong predictions about (66) and (67), and one that is entirely insensitive to questions regarding the hierarchical organisation of syntactic structures and direction in which they are built.

At the end of the day, therefore, it seems that the facts in (62)–(65) provide no support for Phillips' cause. Baltin (2006:745–46) further undermines Phillips' case for top-down left-to-right structure building. He presents an ingenious case of VP fronting and VP ellipsis combined in a single utterance, plus a bound-variable pronoun linked to a quantificational object in the fronted VP, plus the threat of 'infinite regress' under Antecedent-Contained Deletion — a threat that, by the logic of Phillips' reasoning, cannot be averted by attaching the container of the ellipsis site high.

(73) [$_{VP}$ visit *every prisoner*] though I may ___ after *his* laywer does ___, it won't matter

To obtain the bound-variable reading for the pronoun *his* in (73), the adverbial PP *after his laywer does* ___ must, by the logic of Phillips' analysis, be merged in a right-branching structure. But merging the adverbial PP in a right-branching structure delivers an 'infinite regress' problem when we try to resolve the ellipsis. We can straightforwardly avoid 'infinite regress' if we merge the adverbial PP in a left-branching structure, adjoined to the right of the VP. Then, of course, the question is how to get *every prisoner* to c-command the bound-variable pronoun. But recall that quantifiers must, independently, undergo QR to a fairly high position in the clause — minimally *v*P. On the assumption that material adjoined to XP can c-command out of XP, up to the next maximal projection above XP, it should be possible for *every prisoner* to c-command *his* even if the *after*-PP (a temporal adverbial) should turn out to be adjoined to the projection of T. There is, once again, no paradox here.

2.4.1.3 Conclusion

The discussion of 'Pesetsky paradoxes' in section 2.4.1.2 has come up empty-handed when it comes to support for the top-down approach to syntactic structure building. Of the two main syntactic ingredients of Phillips' (1996, 2003) plea for the top-down, left-to-right approach, the binding-theoretic argument turns out to be baseless. In section 2.4.1.1 we had already had occasion to call into question the argument based on temporary constituency as well. There are no temporary constituents that can be destroyed by the addition of an argument later in the derivation; and the kind of destruction of temporary constituency Phillips draws upon in the account of the VP-fronting data with stranding of a container of a bound element linked to a binder in the fronted VP turns out not to be real either.

The literature has invested a considerable amount of energy into 'Pesetsky paradoxes', and has given us fascinating vistas on them, the properties of syntactic structures, and the direction in which they are built. It is disappointing news that they do not, in the end, support the top-down approach to syntactic structure building. But sorting through Phillips' discussion has been an edifying exercise: ascertaining that, initial appearances notwithstanding, apparent conflicts between binding and constituency do not explicitly support a top-down syntactic model has narrowed down the range of possible evidence for the directionality of structure building, and has contributed to our conscious awareness of what does and does not count as evidence for one approach over the other.

In the context of this book, the discussion of 'Pesetsky paradoxes' and 'cascade structures' has another beneficial effect as well: it will prepare us for the existence of genuine cascades that do not have unwanted constituency effects. We will discover such cascades in chapter 5, in the discussion of the syntax of *convince*-type clausal complementation constructions such as *they convinced us that cascade structures really exist* (see section 5.3.14).

In the next section, I will turn my attention to an entirely different kind of argument for a top-down approach, one that will form the spine of the rest of this monograph: the syntax of filler–gap dependencies.

2.4.2 Filler–gap dependencies: Chesi (2007, 2015)

Fong (2005) and Chesi (2007, 2015) have recently pioneered minimalist-inspired approaches to filler–gap dependencies taking a top-down approach and drawing on memory buffers or 'last in, first out' (LIFO) pushdown stacks (Wanner & Maratsos 1978). Their approaches converge on similar conclusions, though their foci are very different, with Fong concentrating more on the parser (i.e., the sentence-processing mechanism) than does Chesi, who stresses that 'despite the similarity with a parsing algorithm, this is not a parsing strategy' (2007:86). Chesi's overall objective is to show 'that while there is no need to conclude that "the parser is the grammar" (in the sense of Phillips 1996:255), nevertheless a top-down oriented derivation, which creates phrase structure from left to right, is compatible with the Minimalist mode of inquiry, and is more processing friendly than the bottom-up approach that is currently standard in Minimalist analyses' (Chesi 2015:71). Because Chesi's goals are more directly aligned with mine in this work than are Fong's, I will set Fong (2005) aside here and concentrate on Chesi's work, with particular emphasis on Chesi (2007), which presents the most explicit discussion of the workings of his syntactic mechanisms.

Chesi's (2007) syntactic goals are *(a)* to make intermediate steps in the formation of long-distance dependencies fall out from the system more naturally, *(b)* to account for strong island effects (subject islands, adjunct islands), and *(c)* to account for parasitic gap constructions. In the following subsections, I will examine whether his top-down left-to-right structure building approach succeeds in delivering a simple perspective on *(a)–(c)*, and improves upon the standard minimalist bottom-up approach.

2.4.2.1 Successive cyclicity

On standard assumptions, syntactic dependencies going beyond the bounds of a single local domain for syntactic computation have to be cut up into several local steps. Chesi reimages the familiar 'successive-cyclic' approach to long-distance dependencies from a top-down perspective, and seeks to derive successive cyclicity from his top-down model.

Chesi explicitly assumes that once the structure-building process hits upon the boundary of an opaque domain (a phase, defined by Chesi as the extended projection of a lexical head, very much as in this book), the memory buffer associated with the higher phase is deposited in that of that domain. For a sentence like *what do you think that he ate?*, Chesi argues that the *wh*-constituent *what* is placed in the memory buffer for the matrix clause as soon as it is encountered because, being merged in a θ′-position, it is unselected in its surface position and therefore has to be linked eventually to its selector. Subsequently, according to Chesi (2007:87), the following things happen:

> The *wh*-phrase in the memory buffer M1 is discharged in the memory buffer of the selected-sequential phase [i.e., the subordinate CP; MdD] ... This memory buffer transmission minimally employs ... re-merge of the content of the memory buffer of P1 in the left periphery of the complement CP ... then, since this position is unselected, the *wh*-phrase is re-stored in the local memory buffer of [the lower clause]. As a result, this "inheritance" mechanism leaves an intermediate copy/trace in the edge of the complement CP phase.

This is supposed to derive successive cyclicity 'without using any formal/edge/teleological feature'. But as Chesi himself points out in a footnote, the assumption that a copy/trace is left behind on the edge of the lower phase 'is not strictly necessary for the algorithm to work' (Chesi 2007:87, fn. 21). He decides to assume it because 'it seems fairly natural'. But the (presumed) requirement of successive cyclicity obviously cannot be formally derived from opinions regarding what is fairly natural. It is not actually the case that Chesi derives the effects ascribed in the literature to successive-cyclic dependencies: he cannot guarantee the presence of a copy/trace on the edge of the lower CP phase; nor does he manage to make the assumption that discharge of *what* in the memory buffer of the lower CP phase requires remerge of *what* in the left periphery of the lower CP follow from anything in his system. Nothing would appear to go wrong if material from the memory buffer of the higher clause could be transmitted to the memory buffer of the lower clause *directly*, instead of being projected onto the edge of the lower clause first and then being transferred to the memory buffer of the lower clause: either way, the *wh*-constituent will be the first element to enter the memory buffer of the lower clause, and should therefore be the last to leave the buffer, as desired.

As a matter of fact, if we do assume that material from the matrix clause's memory buffer needs to be projected onto the edge of the subordinate clause before being transferred to the memory buffer of that clause, it seems that we run into serious empirical trouble in the case of multiple long-distance dependencies.[42] I will illustrate the problem on the basis of the properties of multiple *wh*-fronting in the Slavic languages of the Balkans. It has been known since Rudin's (1988) seminal work on the topic that there is a basic two-way split within the family of multiple *wh*-fronting languages: some (such as Bulgarian) show systematic superiority effects, and arguably front all *wh*-constituents to CP; others (such as Serbo-Croatian) seem to freely allow violations of superiority, at least in single-clause contexts such as (74).

(74)		*Bulgarian*	*Serbo-Croatian*
	a.	koj kogo vižda?	ko koga vidi?
		who whom sees	who whom sees
	b.	*kogo koj vižda?	koga ko vidi?
		whom who sees	whom who sees

42 I thank Márton Kucsera for raising this point.

Importantly, however, in cases of long-distance multiple *wh*-fronting, the difference between Bulgarian and Serbo-Croatian comes out in the wash: now even Serbo-Croatian respects superiority. The example in (75) (Bošković 1997, 1998, 2002) illustrates this for a biclausal case;[43] the facts remain the same in long-distance multiple *wh*-fronting constructions spanning more clause boundaries: systematically, the *wh*-phrase that is structurally superior to the other *wh*-constituent within its own clause must be first in the string in all the Balkan Slavic multiple *wh*-fronting languages whenever the *wh*-dependencies span multiple clauses.

(75) a. ?ko si koga tvrdio da je istukao? (Serbo-Croatian)
who are whom claim.2SG that is beaten
'who did you claim beat whom?'
b. ?*koga si ko tvrdio da je istukao?
whom are who claim.2SG that is beaten

Superiority effects and the varying sensitivity to such effects in single-clause environments are things that a theory of filler–gap dependencies employing memory buffers such as Chesi's can well account for in principle. I will illustrate this later in the book; the details do not matter for now. What concerns us here is whether the consistent superiority effects in long-distance multiple *wh*-fronting constructions can be accounted for in Chesi's system, in which at each subordinate clause boundary the content of the matrix memory buffer must be re-merged at the left periphery of the complement-CP, and be restored in the memory buffer of the lower clause, all in a 'last in, first out' fashion. Consider the schematic representation of the problem in (76):

(76) a. $[_{CP1}\ filler_1\ filler_2$ ...
b. $[_{CP1}\ filler_1\ filler_2$ ... $[_{CP2}\ filler_2\ filler_1$ ...
c. $[_{CP1}\ filler_1\ filler_2$ ... $[_{CP2}\ filler_2\ filler_1$ ... $[_{CP3}\ filler_1\ filler_2$ ...
d. $[_{CP1}\ filler_1\ filler_2$ ... $[_{CP2}\ filler_2\ filler_1$ ... $[_{CP3}\ filler_1\ filler_2$ ... $[_{CP4}\ filler_2\ filler_1$ ...

In the highest clause (CP1), the fillers are stored in the memory buffer of that clause. When we reach the left boundary of CP2, this memory buffer is emptied into the edge of this clause, in LIFO fashion, as depicted in (76b): $filler_2$ is knocked off the stack first, followed by $filler_1$. This is expected to *reverse* the superiority effect seen in the single clause in (76a). If the *wh*'s do not find gaps to bind in CP2, they are restored in the memory buffer of this clause. At the next clause down (CP3), we then repeat the process we went through in (76b) and expect the superiority effect to flip back to the situation in CP1. If there still are no gaps to bind for the two *wh*'s, we do the whole thing again at CP4, reverting back to what we had at CP2. The constant remerging and restoring of fillers in long-distance dependencies thus seems to predict a constant swing of the pendulum back and forth, from clause to clause. No such swings are attested, however: while there is variation with respect to superiority among the Slavic multiple *wh*-fronting languages in a single clause, all these languages show consistent and stable superiority effects in multiclausal environments in which long-distance dependencies are established by each of the fronted *wh*-constituents.

43 I will not illustrate here the fact — unremarkable, in light of the left-hand examples in (74) — that the Bulgarian equivalent of (74) also obeys superiority.

What we have seen in this section is that Chesi's (2007) theory can derive successive-cyclic dependencies only if transmission of left-over material in the memory buffer of a higher clause straight into the memory buffer of the lower clause is impossible. But nothing in the theory guarantees that material from the memory buffer of a higher clause can only end up on the memory buffer of a subordinate clause by first being remerged at the edge of the lower clause and then being restored. In fact, if we did assume this constant remerging *cum* restoring, the LIFO nature of the process would seem to make quite the wrong predictions about long-distance dependencies in multiple fronting constructions. The latter would seem more harmful than failure to derive successive cyclicity; I already pointed out in section 2.3.4 that Chomsky's own bottom-up 'minimalist' approach does not, either.

I conclude that it is unlikely that a top-down approach to filler–gap dependencies exploiting memory buffers can derive successive cyclicity — but that this may not be a problem for the top-down theory at all. I will come back to the matter in chapters 3 and 5.

2.4.2.2 Strong islands

Chesi's (2007) account of strong islands amounts to a particular reformulation of Huang's (1982) CED. The active ingredient is the hypothesis in (77).[44]

(77) 'items in the memory buffer, at the end of the phase, can be transferred only to the memory buffer of the last selected phase' (Chesi 2007:90)

Subjects, in languages such as English, are *per definitionem* never selected phases in Chesi's theory:[45] they are what he calls 'nested phases', and by hypothesis, nested phases cannot be the beneficiaries of the transfer of the memory buffer of a higher phase to a lower one — recall from section 2.4.2.1 that an element placed on hold in a higher phase can only be 'discharged in the memory buffer of the selected-sequential phase' (Chesi 2007:87).

Adjuncts, like subjects, are usually nested phases, and therefore strong islands as well. But we know that not all constituents that must, by standard constituency tests, be classified as adjuncts are opaque for the establishment of filler–gap dependencies across their boundaries. Thus, consider the examples in (78). The former exhibits the familar strong island effect. But the latter (discussed at length in Truswell 2011), does not.

(78) a. *what did John drive Mary crazy [before reading *ec*]?
b. what did John drive Mary crazy [whistling *ec*]?

Chesi is aware of these facts, and tentatively suggests that the bracketed constituent in (78b) and in similar such examples of transparent adjuncts are selected by the verb, and occur in the complement-of-V position of a 'Larsonian' VP. Since the bracketed constituent in (78b) is assumed to be a phase selected by a lexical category (i.e., not a nested phase), it allows the transfer of the memory buffer of the *v*P phase onto itself, making it possible for the filler *what* to be traced back to its gap, *ec*, inside the adjunct.

44 For 'selected phase', Chesi also uses the expressions 'sequential phase' and 'selected-sequential phase'. Note that the formulation of the hypothesis quoted in the main text suggests that direct transfer of material from one memory buffer to another is actually possible — this is relevant in connection with the discussion at the end of the previous subsection regarding Chesi's attempt to derive successive cyclicity.

45 Chesi (2007:92) talks briefly about variation with respect to the opacity of subjects, referring to a proposal by Choi & Yoon (2006).

The idea that transparent adjuncts are selected phrases will come to haunt Chesi, however, in his desire to account for parasitic gap constructions, his third main objective.

2.4.2.3 Parasitic gaps

Key in Chesi's account of parasitic gap constructions is the so-called parasitic use of the content of the memory buffer:

(79) 'a nested phase can use elements in the memory buffer of the superordinate phases but only parasitically, that is, without removing the used elements from the originating memory buffer' (Chesi 2007:100)

Because Chesi (2007:100) characterises the postulation of (79) as an act of 'add[ing] an extra possibility to our memory buffer devices', this immediately starts his account down a rocky road: since its sole purpose appears to be to ensure that parasitic gap constructions can be accommodated by the model, (79) is tantamount to a *pg*-specific device (similar in this respect to Chomsky's 1986 dedicated '0-subjacency' condition on chain composition in parasitic gap constructions). Parasitic gap constructions should fall out without further ado from the UG toolkit; there ought to be no *pg*-specific devices in the system.

But let us indulge, and consider how (79) derives the familiar parasitic gap constructions. It does so straightforwardly: a filler stored in the memory buffer of a particular phase is parasitically associated with a nested phase (a subject in (80a), and an adjunct in (80b)) before being discharged in its selected position.

(80) a. who do [close friends of *pg*] admire *ec*?
b. which article did you file *ec* [before you could peruse *pg*]?

The analysis also explains the well-known 'anti-c-command condition' on parasitic gap constructions, illustrated by such ungrammatical sentences as those in (81).

(81) a. *who *ec* is admired by [close friends of *pg*]?
b. *which article *ec* was filed [before you could peruse *pg*]?

Discharge of an element in the memory buffer of the superordinate phase empties the memory buffer and therefore makes subsequent parasitic use of the element impossible. So the simple cause of ungrammaticality in the examples in (81) is the fact that the fillers are discharged in the positions marked *ec*, which are hierarchically higher than, and therefore reached by the structure-building process before, the bracketed constituents harbouring the parasitic gaps (marked *pg*).[46]

46 Chesi (2007) himself gives (ia) as his chosen example of the anti-c-command condition: the indirect object (on standard assumptions) c-commands the direct object; the direct object is unable to harbour a *pg* licensed by a moved indirect object. (I modified Chesi's own example slightly by adding *wanted to* to make it pragmatically more plausible.) He says in his fn. 36 that anti-c-command leads to the expectation that (ia) should become grammatical once converted into a double-object construction, given that we know that the direct object of the DOC does not c-command the indirect object. It does not seem to be the case, however, that (ib) is significantly better than (ia).

(i) a. *which book did John give *ec* to [a child who wanted to read *pg*]?
b. *which book did you give [a child who wanted to read *pg*] *ec*?

(*continued overleaf...*)

When we combine this account of anti-c-command with Chesi's (2007) assumptions about the nature of transparent adjuncts as 'selected adjuncts', we end up making the wrong predictions about licensing a parasitic gap inside such adjuncts. While the *before*-phrase in (78a) is nested phase, which should unproblematically harbour a parasitic gap (as indeed it can: recall (80b)), Chesi assumes that in (78b) we are dealing with a selected adjunct, located hierarchically lower than any object in the VP: the selected adjunct is the V-head's inner complement, in a Larsonian VP-shell structure. Because of the selected adjunct's extremely low location, any filler on hold in the memory buffer of the matrix phase will have been discharged in its gap position, and therefore removed from the memory buffer, before we ever get to the adjunct. Parasitic use of the memory buffer of the matrix phase in the selected adjunct will hence be impossible. Treating selected adjuncts as inner complements of V thus leads to the prediction that it should be impossible to include a parasitic gap inside a selected adjunct — precisely because no 'parasitic use' of the memory buffer should be possible. But this prediction is false: we can easily turn sentences like (78b) into parasitic gap constructions, as shown in the b–examples below. (The a–sentences are provided to indicate that we are indeed dealing with transparent, 'selected' adverbials here.)

(82) a. who did John drive Mary crazy [talking to *pg*]?
b. who did John drive *ec* crazy [talking to *pg*]?

(83) a. which book did John rewrite his thesis [after reading *ec*]?
b. which book did John rewrite *ec* [after reading *pg*]?

So it seems ill-advised to treat transparent adjuncts as 'selected' constituents that define a selected-sequential phase — at least not all the time: for if they were sequential phases all the time, they ought to be unable to accommodate parasitic gaps. If they are never sequential phases, then they must be nested phases — which should make them opaque by Chesi's logic, which would lose us the account of the contrast between (78a,b). If they are sometimes sequential and sometimes nested phases, the question that arises is how we can tell (other than by inspection of the distribution of gaps) whether a particular non-argument is merged inside or outside the main verb's phase.

Before closing this discussion of Chesi's top-down approach to filler–gap dependencies, I would like to make a brief excursion and relate the parasitic gap construction to the *wh*-scope marking construction, which featured in the discussion in section 2.3 (recall (30) and (31), above). This is the topic of section 2.4.2.4. Readers not interested in this excursion may move on straight to the conclusion, in section 2.4.2.5, and to the chapter's general conclusion, in §2.5.

We may be dealing here with a different kind of restriction (whose identity may not be known at this time) that is responsible for the deviance of both examples in (i). At any rate, it seems unlikely that Chesi's approach to parasitic gap constructions, couched in terms of (79), could have anything to say about (ia,b): the Goal of ditransitive constructions probably constitutes a selected phase rather than a nested one.

It may be that (ib) gets better when *which book* is *sub*extracted from the Theme, as in (iib). Under the same conditions, it seems that (ia) then gets better, too: see (iia). The latter goes along with the classic anti-c-command line: while *ec* c-commands *pg* in (ia), it does not in (iia). But Chesi's recasting of the anti-c-command condition as in (79) will be unable to handle a contrast between (ia) and (iia) (if indeed such a contrast exists): in both cases, if the first object is a selected phase, the memory buffer of the matrix phase will be empty by the time the bracketed constituent comes in.

(ii) a. which book did John give an advance copy of *ec* to [a child who wanted to read *pg*]?
b. which book did you give [a child who wanted to read *pg*] an advance copy of *ec*?

2.4.2.4 On the relationship between parasitic gaps and wh-*scope marking*

A top-down approach to syntactic structure building presents an interesting connection between the parasitic gap and *wh*-scope marking constructions — a connection that I hope future research will find supporting evidence for.

If, as (77) says, 'items in the memory buffer, at the end of the phase, can be transferred only to the memory buffer of the last selected phase' (Chesi 2007:90), no filler outside a category that is not the last selected phase should be allowed to establish a dependency with a trace inside it. If the structure built up to the discovery of the opaque category happens to contain a filler that wants to forge a link with something inside the opaque domain, the derivation is terminated. But if the filler in question can bind a trace in the matrix domain and this filler–gap dependency is linked in its entirety to a separate filler–gap dependency wholly inside the opaque domain, the result should land in its feet.

One context in which a filler–gap dependency inside an opaque domain is linked in its entirety to a filler–gap dependency in the matrix clause we have just encountered: the parasitic gap construction. In (80a,b) (repeated in an updated form in (84)), the bracketed constituents are opaque domains. The operator in sentence-initial position hence cannot itself bind a gap inside the bracketed constituents in (84a,b). But the operator–variable dependencies established *within* these bracketed constituents can be linked up to the operator–variable dependencies established in the matrix clauses (provided that 'connectedness', *à la* Kayne 1984, holds and the 'anti-c-command condition' is satisfied).

(84) a. who do [*Op* close friends of *pg*] admire *ec*?
b. which article did you file *ec* [*Op* before you could peruse *pg*]?

Another context in which two filler–gap dependencies in separate domains, the lower one opaque to the higher one, are linked presents itself in the so-called *wh*-scope marking construction, which Hungarian (85) (which repeats (31a)) illustrates.

(85) *mit* hisz János hogy *ki* terhes? (Hungarian)
what believe.3SG.INDEF János that who pregnant
'who does János believe is pregnant?'

In (85), the *wh*-scope marker *mit* is unequivocally declared a dependent of the matrix verb thanks to the Agree relationship between it and the matrix *v* (as reflected in the indefinite/subjective conjugation). Because the matrix *v* can establish an Agree relation with only one object at a time, the complement-CP introduced by *hogy* is not an Agree goal, and hence, by (47) (repeated below, from section 2.3.2.1), automatically gets declared a opaque domain. It is this entire domain that serves as the associate of the *wh*-scope marker in the higher clause,(y the logic of the so-called 'indirect dependency' approach to *wh*-scope marking constructions of Dayal (1994), Horvath (1997) (see also Bruening 2006, *contra* Bruening 2004).[47]

47 An independent consideration favouring the 'indirect dependency' approach to *wh*-scope marking is based on the resolution of the labelling problem posed by the [XP YP] structure in the subordinate clause. *Wh*-scope marking involves *terminal* movement of a *wh*-phrase to the edge of a subordinate clause, with the 'expletive' *wh*-element marking the scope of the *wh*-phrase. The structure of the embedded clause of a *wh*-scope marking construction would present an insurmountable labelling problem of the type discussed by Chomsky (2013, 2015) if the subordinate clause (an [XP YP] structure) showed no feature-matching Agree relation between the *wh*-phrase and

(47) *opaque domain*
in [α ... π ... [$_{\Delta}$ ... β ...]]], Δ is an *opaque domain* for a relation between α and β iff:
(a) Δ dominates β, and
(b) Δ ≠ a goal γ in an Agree-relation with an asymmetrically c-commanding probe π

The establishment of a link between an operator–variable dependency inside the opaque domain and an operator–variable dependency in a higher structural domain is something that *wh*-scope marking and parasitic gap dependencies share. There is one striking difference, however, between *wh*-scope marking and parasitic gap dependencies: whereas the former involve a secondary operator–variable dependency inside a *complement* clause, the latter resist precisely such a configuration. This resistance is usually captured under the rubric of the 'anti-c-command condition' on parasitic gap constructions — already mentioned in the discussion of Chesi (2007).

The anti-c-command condition on parasitic gap constructions can be derived from Principle C of the Binding Theory (see Mulder & Den Dikken 1992): the parasitic variable would wrongly be A-bound if the variable representing the 'real' gap, with which the parasitic variable is coindexed, were in a c-commanding A-position. If so, we can make immediate sense of the difference between parasitic gap constructions and *wh*-scope marking constructions in this realm when we realise that in *wh*-scope marking constructions the *wh*-dependency in the lower phase is not in a binding relationship with the *wh*-scope marker in the matrix clause: on the 'indirect dependency' approach there is no coindexation of the two *wh*-chains in a *wh*-scope marking construction (unlike in the now widely discarded 'direct dependency' approach; Van Riemsdijk 1983, McDaniel 1989). Because there is no coindexation relation between the two *wh*-chains in a *wh*-scope marking construction, there is no threat of a Principle C violation; hence the anti-c-command condition is not active.[48]

2.4.2.5 Conclusion

Let us return to Chesi (2007, 2015), whose work was the focus of the discussion in this section. Chesi propagates a top-down approach to the construction of syntactic structures and the filler–gap dependencies established within them primarily on the basis of three purely syntactic grounds: *(a)* to derive intermediate steps in the formation of long-distance dependencies ('successive cyclicity'), *(b)* to account for strong island effects (subject islands, adjunct islands), and *(c)* to account for parasitic gap constructions.

C — in other words, if the embedded clause were not interrogative. Labelling considerations thus force an 'indirect dependency' analysis upon the *wh*-scope marking construction: only on such an approach is the subordinate clause treated as a *wh*-interrogative clause associated with the 'expletive' *wh*-element in the matrix.

48 In this regard, *wh*-scope marking constructions are similar to *tough*-movement constructions, on Mulder & Den Dikken's (1992) analysis of the latter, according to which the operator–variable dependency in the infinitival clause of *tough*-movement constructions is not coindexed with the subject of the *tough*-adjective. The parallel between *tough*-movement and *wh*-scope marking constructions can actually be taken further. In *tough*-movement, the constituent harbouring the operator–variable dependency in the lower clause is standardly taken to be predicated of the subject of the *tough*-adjective. For the relationship between the *wh*-scope marker and the subordinate clause harbouring the dependency between the 'real' *wh*-operator and its variable one could likewise argue that it involves predication. This relation is customarily assimilated to that between an expletive and its associate — much as in *I hate it that they did this*. And expletive–associate relations are plausibly thought of as specimens of predication relations: the 'expletive', itself meaningless, gets its semantic content specified by its associate, in an asyndetic specification relationship similar to the one found in sentences like *John met someone important, (namely,) the President of the United States*.

We have seen that Chesi does not deliver on *(a)*, but in this respect his proposal is no worse off than Chomsky's bottom-up work. I will revisit the alleged successive cyclicity of long-distance filler–gap dependencies later in the book (esp. in chapter 4). Chesi's account of strong islands *(b)* is simple and generally satisfactory. But the way in which he seeks to explain the transparency of certain adjuncts (which he calls 'selected adjuncts') does not fit in with his proposed account of parasitic gaps *(c)*, which in itself is little more than a restatement of the problem in terms of memory buffers: it does not provide further insight into the workings of parasitic gap constructions, and is hobbled particularly by the fact that it postulates a *pg*-specific device.

2.5 Conclusion

This scene-setting chapter has reviewed arguments given in the literature for a particular approach to the directionality of structure building and the construction of dependencies within syntactic structures.

It started out its review with a close look at arguments based on the cycle, dating back to Chomsky's (1957) treatment of clausal subordination. Though Chomsky (1965) states that the cycle forces a bottom-up syntactic derivation, and though similar language echoes in more recent work, it turns out upon close scrutiny that cyclic spell-out actually works considerably more efficaciously in a top-down system.

The two extended arguments for top-down derivation given in Phillips' (1997, 2003) work on constituency conflicts are frequently cited not just in the context of the relation between the grammar and the parser (for Phillips, these are one and the same thing) but also in syntax-internal discussions. If valid, Phillips' arguments would be a strong endorsement of top-down syntactic derivation, which is the approach pursued in this book. It is a disappointment, therefore, that we could not conclude in the end that Phillips' arguments hold water.

For filler–gap dependencies, Chesi (2007, 2015) is, to my knowledge, the first to present an explicit and detailed plea for top-down derivation in syntax. The approach taken by Chesi, exploiting a pushdown (or 'last in, first out') stack on which fillers that need to be associated with a position lower in the tree are temporarily placed, is very close in spirit and execution to the one advocated in the following chapters. But the particular analysis presented by Chesi is neither precise enough nor successful in deriving the patterns that need to be accounted for.

Chesi's work as it stands is thus insufficient as a convincing case for a top-down model of syntax. But it is certainly a step in the right direction. In chapter 3, I will take his cue and develop in detail a top-down theory of long-distance filler–gap dependencies and the two different types of locality effects imposed on them: absolute (or 'strong') islands , and intervention (or 'weak') islands. Both absolute islands and intervention effects are shown to fall out naturally from a top-down approach to the construction of filler–gap dependencies, which serves as a selling point for such a model in light of the difficulty that bottom-up approaches have encountered in accounting for them.