CHAPTER 1

INTRODUCTION

From the title and subtitle, readers may surmise that this is a book about three things. It is a book about relative clauses; it is a book about Proto-Indo-European; it is a book about syntax. For readers familiar with these topics, such a book may seem broad in scope, addressing as it does a set of research questions that have occupied several generations of comparative philologists. On the other hand, for all their pertinacity, comparative Indo-European philologists occupy only a relatively small (albeit optimistically increscent) intersection between the larger disciplines of Classics and Linguistics, not to mention the study of the ancient world beyond Greece and Rome. This introduction is aimed primarily at those whose more general interest in these areas may have led them to pick up this book (or click 'download PDF'). I will introduce each of these three themes – relative clauses, Proto-Indo-European, and syntax – and explain how they come together in this study.

Forgoing this introduction, the reader may turn straight to Chapter 2 and Chapter 3, which lay out in detail the methodological and theoretical foundations of the book, respectively; or to Chapters 4–6 for a survey of the relevant Indo-European evidence and the reconstructions themselves; or, most efficiently, straight to Chapter 7 for the conclusions. In the interest of clarity, I avoid using too many abbreviations in the Introduction; those that are adopted in subsequent chapters may be found in the List of Abbreviations, together with glossing conventions.

1.1 Introducing Relative Clauses

That which we call a rose
By any other word would smell as sweet.
So Romeo would, were he not Romeo called,
Retain **that dear perfection which he owes**
Without that title. (Shakespeare, 1597)

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

Relative clauses are everywhere. Though their precise form and function may vary, they are not characteristic of a particular genre or era. My conservative estimate is that this book contains somewhere in the region of 2,000 relative clauses in the body of the text, not including quoted examples – there were four in the first two paragraphs. Nevertheless, defining in precise and universally applicable terms what exactly a relative clause is is remarkably difficult (see §3.1). By way of an introduction, then, the most expedient way to approach relative clauses is to start with an example. Of those highlighted in bold in the excerpts above, let us take the one that is easiest to dissect:[1]

(1.1) *That dear perfection* [Relative clause *which* ***he owes***]

A simple but important fact about relative clauses is that they are **clauses**. In rudimentary terms, a clause is something that has a subject and a verb. Each of the examples quoted above shows these characteristics: in (1.1), it is *he owes*. But what makes these clauses specifically relative is that they provide some more information about an entity – in technical terms we can describe this as **modifying** it. Traditionally, the modified entity is referred to as the **antecedent** of a relative clause, or the **head noun** in more contemporary terminology. In (1.1), the head noun is the noun phrase underlined: *that dear perfection*. The relative clause, then, tells us that *that dear perfection* is something Romeo possesses. The nature of the modification is encoded in the fact that the head noun must fulfil a syntactic role in the relative clause. In (1.1), the head noun refers to the object of the relative clause. We can imagine a counterpart of the form *He owes that dear perfection*, where *that dear perfection* follows the verb in the usual position for objects. When we turn this into a relative clause, the head noun appears outside the relative clause, and there is a gap in its place: *That dear perfection which he owes* ____ .

Relative clauses cannot stand alone: in isolation, (1.1) does not make much sense. This is part of what makes them a prime example of a **subordinate clause** (see §3.1.1). A subordinate clause

[1] *He owes* here can be taken to mean *He has*.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

has to be integrated into a main clause to be felicitous. So, to put (1.1) in context:

(1.2) [Main clause *Romeo would retain that dear perfection* [Relative clause *which he owes*] *without that title*.]

When we integrate a relative clause into the main clause, we notice its most intriguing aspect: the fact that the head noun, while playing a grammatical role in the relative clause, also has a role to play in the main clause. We saw in (1.1) that the head noun *that dear perfection* was the object of the relative clause *he owes*. In (1.2), we see that, together with the relative clause, the whole entity described as *that dear perfection which he owes* is itself the object of the main clause, *Romeo would retain*. In technical terms, the head noun + relative clause combination is a complex nominal: it can fulfil the same function as a normal noun phrase (i.e. it can be a subject or object), but it contains a separate clause within it. Two simple clauses with the same noun phrase would have done the job: *Romeo owes that dear perfection; he would retain that dear perfection without that title*. Using a relative clause combines the two into one complex structure.

I have not yet mentioned the most textbook, telltale sign that we are dealing with a relative clause here: namely, the presence of a **relative pronoun** that introduces it. In English, this typically takes the form *who* or *which*. We see it in (1.1), and elsewhere in the Shakespeare passage. In English, however, we cannot always rely on the presence of the relative pronoun to signal the introduction of a relative clause: indeed, we can grammatically (if not stylistically) do without the relative pronoun in (1.2).

(1.3) *So Romeo would, were he not Romeo called, retain that dear perfection* [Relative clause *he owes*] *without that title.*

The presence or absence of relative pronouns is something that varies cross-linguistically, which is part of why a global definition of a relative clause is so elusive (it cannot be defined by the presence of a relative pronoun, for instance). When it comes to the languages that most concern this book, however, the obligatory presence of a relative pronoun is an exceedingly common feature. Those familiar with an ancient Indo-European language, such as

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

Latin, will find relative clauses immediately recognisable by this token:

(1.4) *O Libane, uti miser est*
O Libanus.VOC.SG how sad.NOM.SG.M be.PRES.3SG
homo ***qui*** *amat*
man.NOM.SG.M REL.NOM.SG.M love.PRES.3SG
'O Libanus, how miserable is a man **who**'s in love'
(Plaut. *Asin*. 616)

(1.5) *sed quis hic est*
but INT.NOM.SG.M this.NOM.SG.M be.PRES.3SG
homo, ***quem*** *ante aedis*
man.NOM.SG.M REL.ACC.SG.M before house.ACC.PL.F
video hoc noctis?
see.PRES.1SG this.ABL.SG.N night.GEN.SG.F
'But who is this man, **whom** I see before the houses at this time of night?' (Plaut. *Amph*. 292)

Equipped with a morphology richer than that of English, we can also identify more readily the syntactic role the head noun plays in the relative clause, making explicit what is implicit in example (1.1). In (1.4), since the modified element *homo* is the subject of the relative verb *amat*, the relative pronoun *qui* appears in the nominative case; in (1.5), since the head noun *homo* is the object of the relative verb, *video*, it appears in the accusative (*quem*). In all cases, the relative pronoun agrees in number and gender with the head noun. The situation in Latin is typical of the way relative clauses are formed in ancient Indo-European languages, and the relative pronoun is central to the construction. This leaves us with much to say about relative pronouns in Indo-European, and they are treated in detail in Chapter 4.

So, why a book on relative clauses? There are two sets of motivations. The first is the importance of relative clauses to understanding a grammatical system. We may rather crudely divide the grammatical elements of a language into two major categories: nominals, which refer to entities; and clauses, which refer to propositions about those entities. As complex nominals, relative clauses at some level allow us to embed a clause within a nominal,

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

or a proposition within an entity. Understanding how a language manages to achieve this feat gets to the heart of how sentences are built, and where the generative power of language lies (I return to this in §1.3 below). All of which is to say, a construction whose nomenclature may be associated with the dustier pages of a reference grammar is in fact far from peripheral in its function, and provides a vital source of evidence for grammatical structure.

Beyond the more general linguistic interest of relative clauses, there is a set of motivations that are particular to the Indo-European language family, arising from the fact that, for all the substantial grammatical diversity among these languages, the relative clause is a remarkably resilient construction. It survives across the various subgroups of Indo-European (see §1.2 below), from the very earliest attestations to the present day. What's more, relative clauses of the Indo-European type (viz. finite clauses introduced by a specialised relative pronoun) are typologically rare outside this language family. This reinforces the notion that where we see similarities among the oldest Indo-European languages, these are most likely not chance resemblances or parallel evolution, but rather something that has been inherited from their shared ancestor. Given the challenges faced when reconstructing an unattested language's syntax, especially for Proto-Indo-European (see §2.1), the possibility that relative clauses in the attested languages could reflect aspects of the syntax of the proto-language directly makes them a point of virtually inevitable interest to the philologist, such that they have received continual scholarly interest for the best part of two centuries. What has changed substantially in that time, which makes an up-to-date treatment timely, is our knowledge of the ancient Indo-European languages and their syntax.

1.2 Introducing Proto-Indo-European

It is probably safe to say there are more books that take Indo-European philology as their subject than those that focus specifically on relative clauses. The reader may peruse the introductory chapters of Anthony (2007), Clackson (2007) or Ringe and Eska (2013) for a more thorough introduction to the discipline than can fit into this work. I discuss the vital concepts of relatedness and the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

Comparative Method in detail in §2.1.1; here I give an outline of the Indo-European language family and the relationships between the languages that feature most prominently in this book.

The Indo-European language family consists of some 400–500 languages. These may be organised into a family tree or **phylogeny**, where closely related languages cluster under a branch or subgroup such as Romance (French, Italian, Spanish etc.) or Indo-Aryan (Hindi, Marathi, Bengali etc.). Winding back the clock on the attested languages as far as possible, we may reduce the number of 'primary' branches of Indo-European to twelve, more or less.[2] These are, in alphabetical order: Albanian, Anatolian, Armenian, Baltic, Celtic, Germanic, Hellenic, Indo-Iranian, Italic, Phrygian, Slavic and Tocharian. All the languages belonging to these families, it is hypothesised, are descended from a single ancestor language: **Proto-Indo-European**. Since Proto-Indo-European is not attested directly, it is impossible to give a precise period of time during which it was spoken, but it is unlikely to have been spoken any later than the fifth millennium BCE.[3] And since it is not attested – there are no surviving Proto-Indo-European texts – to say anything about it at all we have to **reconstruct** it, based on what we see in its attested descendants. A basic sketch of the relationship between these descendant languages, tracing them back to Proto-Indo-European, is given in Figure 1.1.[4]

The phylogeny in Figure 1.1 shows the twelve 'first order' subgroups in capital letters, and is helpfully detailed in terms of the languages listed beneath them. There are two critical aspects of the Indo-European evidence, however, that are not presented in Figure 1.1. The first is the dates of attestation for the earliest evidence for each language. While one might be tempted to take the 'first order' subgroups as occupying a similar place in time, the reality is something quite different. The oldest attestations of any Indo-European language are texts in Hittite, an Anatolian

[2] The precise phylogeny of Indo-European is a matter of considerable debate, and there are various different ways of approaching the task. See *inter alia* Olander (2022).

[3] See most recently Kloekhorst (2023).

[4] There are some quirks in this phylogeny, especially concerning the Celtic languages. Gaulish is listed incorrectly under Brythonic, and Celtiberian is not present at all: see Sims-Williams (2015: 357–61) for informed discussion.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

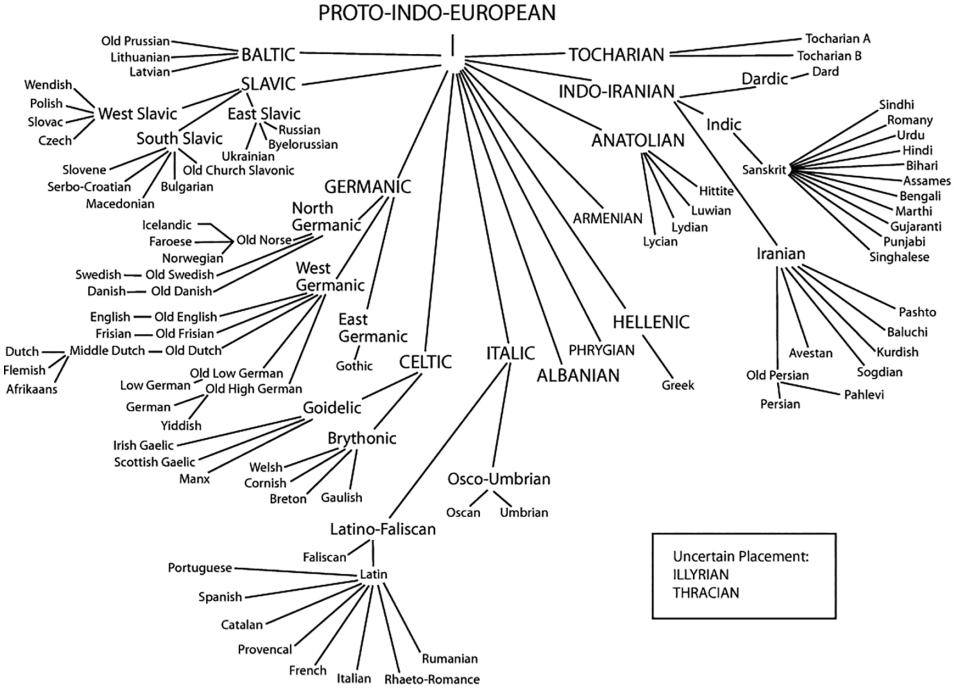


Figure 1.1 The Indo-European language family (Anthony, 2007: 12). Used with permission of Princeton University Press. Permission conveyed through Copyright Clearance Center, Inc.

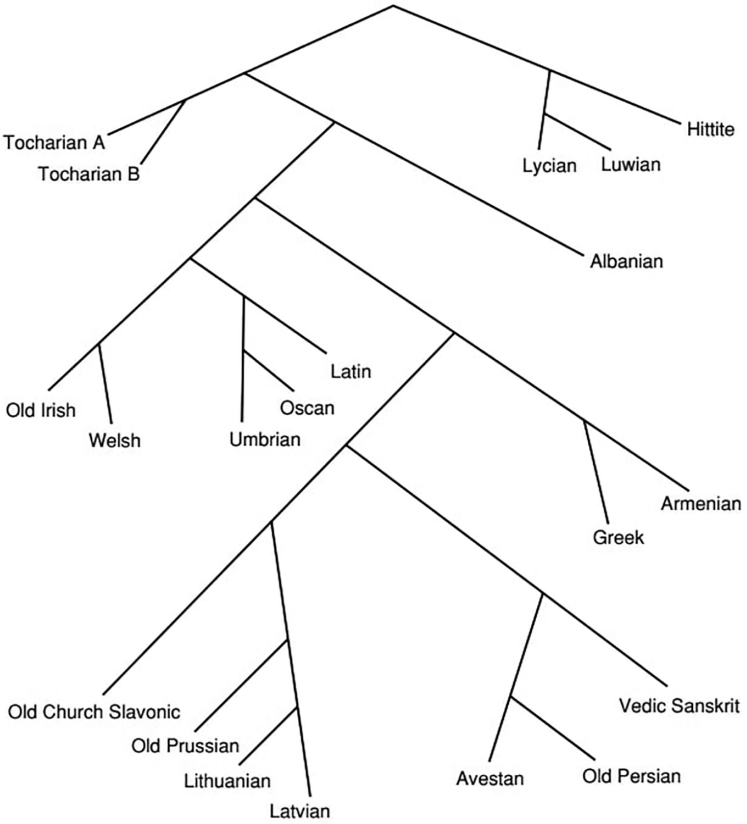


Figure 1.2 The 'Pennsylvania Tree' of Indo-European languages (Ringe *et al.*, 2002: 100). Used with permission of John Wiley & Sons. Permission conveyed through Copyright Clearance Center, Inc.

language attested from c. 1800 BCE; on the other hand, the oldest evidence from the Tocharian languages dates to the fifth century CE; the most archaic Baltic evidence, from Old Prussian, appears as late as the fourteenth century. Another aspect that is hidden in Figure 1.1 is the fact that there is cogent linguistic evidence that subgroups exist beyond the 'first order' branches. Figure 1.2 shows an alternative phylogeny, derived by Ringe *et al.* (2002: 100) on the basis of a variety of phonological and morphological features. The unmarked intermediate nodes along the spine of the tree may be

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

considered 'second order' subgroups, each representing in effect its own unattested proto-language that subsequently separated out into the 'first order' subgroups we see in Figure 1.1.

Trees like this start to tell us a story about the diversification of Indo-European languages. The first to 'split off' are the Anatolian languages (illustrated with Hittite, Lycian and Luwian in Figure 1.2), which lack a substantial number of features that are shared among all other branches. Scholars have come to refer to this as the **Anatolian split**, the first stage in the breaking up of the Indo-European linguistic community, and the *terminus ante quem* for the latest stage of Proto-Indo-European linguistic unity. The next group to split off is Tocharian; this change is much more difficult to date, because of how late the earliest attestations of Tocharian are, but the linguistic distance of Tocharian from the rest of the family means that there is general consensus that it was probably second after Anatolian (see Peyrot, 2019b). After Tocharian, the picture becomes less clear.

Indeed, the precise phylogeny of the Indo-European language family remains a matter of considerable debate. Around the same time as Ringe *et al.* (2002), Gray and Atkinson (2003) applied Bayesian methods[5] to a large data set of cognate vocabulary to assess the most likely genealogical connections between languages, as well as estimating the time when the different branching events occurred. Since then, Bayesian phylogenetic analyses have been pursued by several authors, with differing results. These include Bouckaert *et al.* (2012), Chang *et al.* (2015), Rama (2018), Ritchie and Ho (2019) and, most recently, Heggarty *et al.* (2023), whose phylogeny is reproduced in Figure 1.3. An altogether different approach, the Parameter Comparison Method, eschews cognate vocabulary as a measure of relatedness, focusing instead on syntactic parameters (see, e.g., Guardiano *et al.*, 2020). The number of parameter values shared between two languages allows us to measure the syntactic 'distance' between any two languages; these distance measures then provide the input for the phylogenetic algorithm. While the overall shape of the Indo-European

[5] See Greenhill *et al.* (2020) for an introduction to these methods as applied to linguistic phylogeny.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

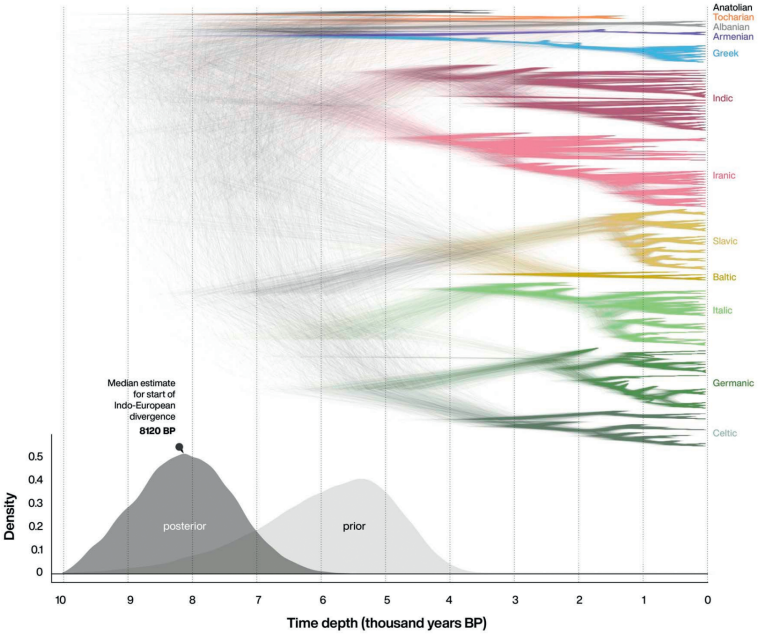


Figure 1.3 The posterior probability distribution of trees for the Indo-European family. BP = 'Before present', with the present taken to be 2000 CE (Heggarty *et al.*, 2023: 4). Used with permission of American Association for the Advancement of Science. Permission conveyed through Copyright Clearance Center, Inc.

language family is maintained, once again the details of subgroupings vary (see Figure 1.4).[6]

The phylogeny of the Indo-European family thus remains an active area of research. For the purposes of this book, however, I only rely on the existence of one 'second order' subgroup: that which includes all the Indo-European languages except Anatolian. In other words, I follow the scholarly consensus in assuming that after the Anatolian split there was a further period of linguistic unity before the splitting-off of Tocharian and then subsequent varieties. I do not use the term Indo-Hittite to refer to the family pre-Anatolian split, as some scholars do, but where necessary – especially in the discussion of Wackernagel's law and the left

[6] In fact, the authors produce a phylogeny that relates to both Indo-European and non-Indo-European languages; I omit the non-Indo-European part here.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

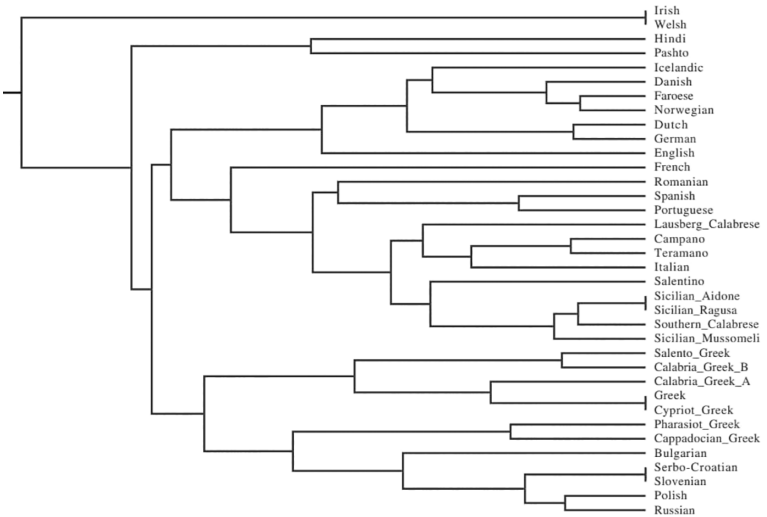


Figure 1.4 Indo-European phylogeny based on parametric comparison (Guardiano *et al.*, 2020: 167). Used with permission of John Wiley & Sons. Permission conveyed through Copyright Clearance Center, Inc.

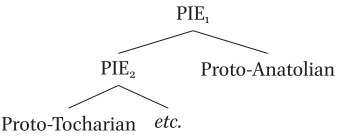


Figure 1.5 The first two Indo-European splits.

periphery in Chapter 5 – I refer to PIE_1 to refer to the period before the Anatolian split and PIE_2 to refer to the period immediately after, before the Tocharian split. Where I write simply PIE, I refer to the period **before** the Anatolian split, i.e. the latest period of comprehensive linguistic unity. This is summarised in Figure 1.5.

A final note is necessary concerning the way I use the term PIE, in comparison to the way it is sometimes used by other authors. I use the term specifically to refer to the latest stage of linguistic unity before the Anatolian split. I work with the assumption that, as with any language, Proto-Indo-European has its own diachrony; it is theoretically possible then, by a process of internal reconstruction (aiming to explain and/or reduce any irregularities left

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

in the system after comparative reconstruction has run its course), to reconstruct some aspect of Proto-Indo-European that predates this latest stage of linguistic unity; we may call this stage 'Pre-Proto-Indo-European', whether we consider this simply to be an earlier form of Proto-Indo-European or a direct ancestor to it. For the purposes of the current work, I will not concern myself with Pre-Proto-Indo-European. It is hard to say with certainty that the Proto-Indo-European we reconstruct using the Comparative Method can be treated holistically as a synchronic system; one can well imagine that certain lexical items we reconstruct predate others within the Proto-Indo-European linguistic system, and the same could just as well be said for syntactic structures of the sort discussed in this book. In a bid to recognise the existence of diachrony in the proto-language without trying to solve it, I explicitly avoid trying to reconcile apparent irregularities or complexities as they pertain to the system of relative clauses, or appealing to synchronic typological universals where the comparative evidence runs out. This has consequences especially for the form of the relative pronoun, discussed in Chapter 4, and the status of embedded relative clauses, discussed in Chapter 6.

1.3 Introducing Syntax

I turn to the third of the three topics that concern this book: syntax. In the etymological sense of the term (from Ancient Greek σύνταξις), syntax is the 'putting together' of linguistic elements (i.e. words, or **lexical items**). Building on this definition, but integrating it into a more contemporary mentalist view of grammar, we could define syntax as the abstract cognitive program that 'puts together' linguistic elements to form sentences. This introduces the notion that syntax is something productive or **generative**: it is a program that takes a given input – we may conceptualise this as items taken from the **lexicon** – and generates a given output, i.e. a sentence. When we analyse the syntax of a sentence, then, we are effectively aiming to surmise what kind of generative program (or **grammar**) could have produced it. To take a concrete example, consider the following:

(1.6) *John loves flowers*.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

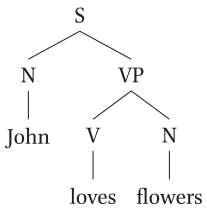
Generally speaking, English exhibits S(ubject) V(erb) O(bject) word order. We see this in (1.6). The subject is *John*, which precedes the verb, *loves*, which in turn precedes the object, *flowers*. Our syntactic analysis of this sentence will thus lead us to postulate some part of the mental grammar that produces SVO sentences. We can do this with a couple of phrase-structure **rules**:[7]

(1.7) S → N VP

(1.8) VP → V N

If we run this two-step program, selecting different Vs and Ns from the lexicon, we can generate every possible SVO sentence in the language. Each of them will have the same structure, which can be represented by a tree diagram such as the following:

(1.9)

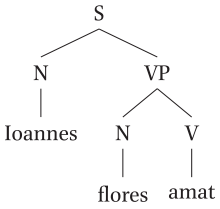


In a language that has SOV word order, all we need to do is switch around the elements in rule (1.8): VP → N V. So, the Latin sentence:

(1.10) *Ioannes* *flores* *amat*
John.NOM.SG.M flower.ACC.PL.M love.PRES.3SG
'John loves flowers'

May be analysed as:

(1.11)



[7] Where S = Sentence, N = Noun Phrase, VP = Verb Phrase.

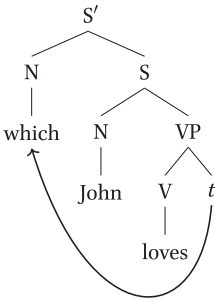
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

Of course, the word order of a language is often more complicated than can be summarised with a singular template such as SVO.[8] Take the following relative clause, a counterpart to (1.6):

(1.12) These are the flowers [$_{\text{Relative clause}}$ which John loves.]

In this example, the relative pronoun (*which*) appears to function as the object of the verb *loves*; yet, rather than appearing after the verb in strict SVO order, *which* appears at the beginning of the relative clause. This order cannot be generated by the two rules posited above, so we need to extend the grammar to account for it. One way of doing that, which unifies the syntax of (1.6) and (1.12) keeping the generalisation that objects follow verbs, is to allow for the **movement** of the object *which* to the beginning of the relative clause; *which* does this because it is a relative pronoun (or ***wh*-word**), whereas *flowers* in (1.6) does not. To represent this movement on a tree diagram:[9]

(1.13) These are the flowers . . .



Movement of *wh*-words to the front of the clause is cross-linguistically common, and referred to as *wh*-movement; we see it in relative clauses in Latin too:

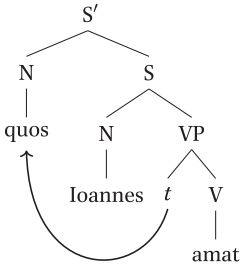
[8] Indeed, adding to this complication is the fact that in a richly inflected language such as Latin, the word order itself does not serve to distinguish between subject and object; that is the job of inflectional morphology. Thus, inflection is also implicated in syntactic rules.

[9] *t* here denotes the trace left behind by the movement of *what*, explaining why the space can't be filled by another element; cf. the unacceptability of something like **These are the flowers which John loves roses*.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

(1.14) *hi* *sunt* *flores*
DEM.NOM.PL.M be.PRES.3SG flower.NOM.PL.M
[Relative clause *quos* *Ioannes*
[Relative clause REL.ACC.PL.M John.NOM.SG.M
amat]
love.PRES.3SG]
'These are the flowers [Relative clause which John loves]'

(1.15) *hi sunt flores* . . .



By combining a sufficient set of rules, such as (1.7) and (1.8), with a sufficient set of possible movement operations, such as *wh*-movement, we can start to map out the content of a language's syntax. This is the basic premise of Chomskyan generative grammar, with origins in Chomsky (1957).[10] The precise status of 'rules' and 'movement' within the grammar have undergone substantial theoretical revision since then, with major overhauls in the 1980s (*Government and Binding Theory*, Chomsky, 1981) and 1990s (the *Minimalist Program*, Chomsky, 1995). The grammatical framework adopted in this book draws on different elements of these latter two iterations of Chomskyan generative grammar (see §2.6 for details). While the basic function of the syntax has been maintained in some form or another, the hypothesised content of the syntax has changed dramatically since Chomsky (1957); there are

[10] Many other grammatical frameworks exist beyond those of Chomsky, notably Construction Grammar (see Hoffmann & Trousedale, 2013), Lexical Functional Grammar (see Dalrymple *et al.*, 2019), Dependency Grammar (see Osborne, 2019), Head-Driven Phrase-Structure Grammar (see Müller *et al.*, 2021), Categorial Grammar (see Morrill, 2010) and so forth. A full-length introduction to syntactic theory would cover these, but in the interest of space I present in detail only the view of syntax that informs the theoretical framework adopted in this book.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

no phrase-structure rules as such, nor is there movement; there is just one operation, MERGE, that takes lexical items and builds them into a hierarchical structure that encodes grammatical information such as subject and object: information we see on the surface as, for instance, orderings of S, V and O or case inflection. The main propositions of this book are grounded within a (largely) Minimalist framework; there is no scope for a detailed introduction here, but there is a plethora of didactic material available, including textbooks such as Radford (2004), Boeckx (2006) and Van Gelderen (2017). I discuss here only the 'big picture' aspects of Chomskyan grammar as they pertain to the specific goals of historical and reconstructive syntax.

The first and most important aspect of this grammatical framework is the understanding that when we talk of the 'syntax of a language', what we are really talking about is 'the set of structure-building operations that underlie the sentences of that language'. In other words, each sentence is taken to be the creative output of an individual's mental grammar. This raises a certain challenge for the Comparative Method, discussed in detail in §2.1. The vocabulary of comparative philology – 'ancestors', 'descendants', 'inheritance', 'relations' – reveals its metaphorical grounding in genealogical biology. The notion of 'inheritance' is particularly important: in the same way that two related organisms may inherit some portion of their DNA from a common ancestor, languages may 'inherit' some portion of their 'linguistic DNA' from their shared ancestor language. Traditionally, this 'linguistic DNA' is found in a language's lexicon. The comparison of individual words or lexical items, inherited via the language acquisition process from one generation to the next, may attest two languages' relatedness in an analogical way to the comparison of genetic data in organisms. We can abstract away from those specific lexical items, compare the languages' sounds (phonology) and word-formation (morphology), because these are also inherited, via the inheritance of lexical items. The challenge of syntax, then, is that sentences are not generally inherited, but rather represent the creative output of an individual's mental grammar each time they produce an utterance. Sentences by themselves, therefore, cannot be the locus of comparison in the same way as inherited lexical items. Rather, it

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

is the structure-building operations or mental grammars of these speakers that are to be compared; the question as to what extent grammars can be inherited in a way analogous to lexical items is something I address in detail while establishing a methodology for syntactic reconstruction in Chapter 2.

The second aspect of this framework, which may be equally important for the purposes of linguistic theory, but less so for the purposes of this book, is that Chomskyan grammar is a nativist theory. This means that it works with the assumption that whatever aspect of the human mind allows us to acquire language, it is (a) genetically inherent and (b) universal to all human beings. These assumptions are brought together under the umbrella of Universal Grammar. The basic idea is that, since all humans have the innate capacity to acquire language, and they do this quickly and effortlessly as infants, they must have some form of biological guidance that constrains the set of possible rules (or grammars) they have to consider.[11] One way of modelling this 'guidance' is found in the framework of Principles and Parameters (Chomsky, 1981): the notion that all languages share some underlying *principles*, and all cross-linguistic variation can be codified in a finite set of *parameters*. The theory of Principles and Parameters, and indeed Universal Grammar more generally, has received much criticism over the years.[12] A major difficulty is that questions of nativism are hard to prove, one way or the other. This means one of the fiercest battlegrounds concerns what the null hypothesis is: whether we should expect all languages to share certain features, beyond the trivialities of human physiology, or whether we should not. Despite my adoption of Minimalist analysis, I plant myself firmly on the fence when it comes to the larger questions of nativism and Universal Grammar; such questions are not those I set out to answer in this book. For our purposes, the value of a universalist framework lies in its comparative potential; the adoption of certain universalist assumptions, all other things being equal, allows us to isolate points of grammatical correspondence that are relevant for the purposes of reconstruction. To treat the syntax of each language

[11] This is the argument from the 'poverty of the stimulus', on which see Lasnik and Lidz (2017).

[12] See, among many others, Evans and Levinson (2009).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

as a *tabula rasa*, ignoring everything we know of other languages, renders the task of syntactic comparison, if not insurmountable, something that is heavily dependent on impressionistic observations of similarity, which are insufficient for the purposes of the Comparative Method.

1.4 Summing Up

With introductions made, I pull together these three strands of inquiry to motivate this book. The structure of relative clauses falls definitionally within the domain of syntax. The different aspects of the structure, from its overarching status as a complex nominal to the details of case behaviour and position of the relative pronoun, are accounted for by the different structure-building processes involved in the syntax. While other aspects of relative clauses, in particular the precise form of the relative pronoun, draw on the Comparative Method as applied to phonology and morphology, a full account of relative clause structure in any language will necessarily be a syntactic one. As discussed in §1.1, relative clauses are a well-established feature of Indo-European languages and, as such, are a prime candidate for reconstruction in some way or another for Proto-Indo-European. This becomes all the more pressing in light of our understanding of syntax as a generative program introduced in §1.3; one way to address the challenge of syntactic reconstruction is to restrict the locus of comparison to a set of syntactic phenomena that are primed for comparative reconstruction by the presence of a demonstrably inherited lexical item: the relative pronoun, whether of the form *$k^wí$/*$k^wó$- or *$yó$- (on which see Chapter 4).

I save a fuller discussion of the implications of this study for the Conclusions. Here it is sufficient to list just two. First: a comprehensive reconstruction of relative clause structure significantly advances our knowledge of Proto-Indo-European. Much progress has been made in reconstructing various aspects of the language in the last two centuries, but major advances have typically been limited to phonology and morphology. We know comparatively little of its syntax: an analysis of relative clause syntax presents a window into a grammatical system that remains largely opaque.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001

Second: that syntactic reconstruction can indeed be undertaken successfully, something that has been the subject of fierce debate in the last fifty years, suggests that this methodology can be extended in principle to other linguistic families. In other words, leaving behind the Indo-European languages altogether, Minimalist reconstruction as executed in this book may be applied to other reconstructed languages across the globe, whose material remnants are lost but whose linguistic legacy lives on in their descendants.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 14:44:39, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009664196.001