

Introduction

The goal of this volume is to explore interactions between syntactic structure and discourse structure, through a case study of patterns of extraction from coordinate structures. It is not a typical survey monograph because the primary empirical focus is a fairly obscure and recalcitrant corner of locality theory. However, this obscure and recalcitrant corner serves us well as a microcosm, and allows us to illustrate fundamental methodological points about the interactions between syntax and discourse.

As is typical for a survey, there is not much original research in what follows. Nevertheless, we believe that this is the most complete account of extraction from coordinate structures to date. This is a consequence of the theoretical breadth of the survey: extraction from coordinate structures is, at first blush, a syntactic matter, but the survey ranges far beyond syntax, and this breadth raises theoretical and empirical questions across syntax, semantics, pragmatics, and discourse structure. A complete survey of extraction from coordinate structure must pay attention to all of these domains, and their interactions.

It will quickly become clear that we are not aiming to promote a single analysis. Instead, we want to motivate reasonable hypotheses which allow us to reason deductively from empirical facts to theoretical conclusions. The theoretical conclusions are likely to be significant in scope: we aim to show that this empirical area has the potential to enable us to discriminate between current syntactic theories, and to inform work on the interfaces between the domains just listed. However, in many cases, the necessary empirical work has not yet been done, and too much of the literature revolves around the same handful of examples, mainly in English. We hope that this survey will inspire further work on extraction from coordinate structures, particularly in understudied languages, and provide a guide to how to tease out the theoretical implications of empirical findings.

In the rest of this introduction, we introduce some overarching issues concerning syntax–discourse interactions, and then motivate our choice of case study.

1.1 Syntax, semantics, and discourse

Although it is now common to see the term ‘syntax–discourse interface’ in the literature (see, for instance, Avrutin 1999, Burkhardt 2005, and Erteschik-Shir 2007), the term is surprising, because such an interface seems to straddle the divide between competence and performance. Syntax, as part of competence grammar, should interface with other parts of competence grammar (morphophonology and semantics), while discourse effects are usually analyzed, with pragmatics, as performance phenomena, concerned with language use in context.

This means that we should be alert to the possibility of fundamental differences between the syntax–discourse interface and, say, the syntax–phonology interface, because the latter, as a mapping between two declarative ‘competence’ representations, is of a different nature from the former, with one foot in competence and one foot in performance.

There are two overly simple responses to this challenge, which we will not pursue seriously. The first is to say that empirical phenomena involving the ‘syntax–discourse interface’ are all performance phenomena. It will be hard to take this seriously because the empirical facts laid out in Chapters 3 and 4 in particular are too intimately related to locality phenomena which are central to the study of syntax in competence grammar.

The second simple response is to say that the syntax–discourse interface is somehow prefigured in syntactic structure, through syntactic counterparts of discourse relations. An example of this approach would be the cartographic syntax of Rizzi (1997), where syntactic [TOPIC] and [FOCUS] features derive syntactic restrictions (for instance on word order) with correlates in discourse structure. The gist of the problem with this approach, regardless of its general merits, is that syntactic structure demonstrably is not discourse structure, and it will turn out that we need to refer to both, and find a vocabulary for describing the interactions between the two structures.

This survey instead follows a separate path, one which is possible only because the landscape around representations of meaning above the sentence level has changed. Research into cross-sentential binding phenomena by Kamp (1981) and Heim (1982) led to the development of formal semantic theories in which the status of semantic correlates of the sentence is diminished, and sentence meanings are compositionally integrated into larger representations of discourse-level information. In short, this body of work motivates a redrawing of the line between competence and performance, and a conceptualization of part of the syntax–discourse interface as belonging within competence grammar, as a component of the syntax–semantics interface.

In Chapter 6, we will focus on a particular extension of Kamp’s Discourse Representation Theory (DRT), namely SDRT, or Segmented Discourse Representation Theory (Asher and Lascarides 2003), which embodies rich hypotheses about discourse structure. As we shall see, representations of discourse structure share certain properties with representations of syntactic structure (e.g. both make use of a relation of hierarchical embedding, or **subordination**), but are incommensurate in various ways (for instance, discourse subordination relations do not always map onto syntactic subordination relations).

This approach to the syntax–discourse interface gives rise to several foundational questions, including the following:

- Are certain theories of syntactic structure more congenial to theories of discourse structure, and vice versa?
- Which structures represent which empirical phenomena?

Questions like these are familiar from other work on linguistic interfaces. Linguistic theory is a long way from providing definitive answers to them, so the best way to make progress is through case studies such as this one, chosen with one eye on empirical tractability and one eye on theoretical richness across subdisciplines. The latter consideration is important because progress in linguistics can become encapsulated within subdisciplinary communities, such as ‘syntactic theory’ or ‘discourse semantics’. Working on an interface

like this one necessarily involves drawing out connections between those subdisciplines. Perhaps unsurprisingly, the empirical discoveries that those subdisciplines have made, and the conclusions that researchers have drawn on the basis of those discoveries, don't always translate across disciplinary boundaries. Articulating what a realistic synthesis looks like, therefore, is a difficult challenge. This volume illustrates that challenge, as well as the virtues of meeting the challenge.

Instead of tackling questions like the above head on, we will identify a series of choice points in the course of this survey (collated in Appendix for easy reference). These choice points are interrelated, and range across syntax and discourse. Ultimately, they are all empirical questions, but in our opinion we do not yet have the data to resolve them. We hope that the choice points set an agenda for future research on extraction from coordinate structures, and on the syntax–discourse interface, by identifying the core questions, the relationships between them, and ways in which answers to these questions might be found.

1.2 Extraction from coordinate structures

Our case study concerns correlations between the interpretation of coordinate structures and patterns of extraction from those structures. Extraction from coordinate structures is limited. In fact, Ross (1967) proposed the **Coordinate Structure Constraint**, or CSC, which prohibits extraction of conjuncts, and extraction out of conjuncts. The effects of these two parts of the CSC are illustrated in (1) and (2), respectively.

- (1) a. Who did you meet __?
b. *Who did you meet [__ and Sue]?
- (2) a. What did you eat __?
b. *What did you [[eat __] and [drink water]]?

There are several known classes of counterexample to the CSC. Some can be characterized in purely syntactic terms. For instance, we will see in Chapter 3 that some languages do, in fact, allow extraction of initial conjuncts, of precisely the sort which (1b) shows to be ungrammatical in English. These examples are interesting in their own right, but not particularly illuminating for the study of the syntax–discourse interface, so we will draw a line around them and leave them to one side, for the attention of specialists in 'pure' comparative syntax.

We focus instead on patterns of extraction from one or more conjuncts, each apparently correlated with a particular interpretation of the coordinate structure. The first, and most discussed, pattern is **Across-The-Board movement**, or ATB extraction, also first observed by Ross. In ATB extraction, a single moved phrase corresponds to a gap in each conjunct. Often, but not always, ATB extractions correspond to **parallel** or **contrastive** interpretations. For instance, in (3), *Kim* and *Sally* contrast, as do *enjoy* and *hate*.

- (3) What did [[Kim enjoy __] but [Sally hate __]]?

Three more correlations were established by Lakoff (1986), building on observations amassed by researchers over the previous twenty years. In (4) (Lakoff's 'Type A'), the

coordinate structure is interpreted as a **narration** and gaps are located in one or more typically noninitial conjuncts.

(4) What did you [[go to the store] and [buy __]]?

In (5) (Lakoff's 'Type B'), the interpretation involves a **violated expectation**, and the gap is in the initial conjunct only.

(5) How much can you [[drink __] and [still stay sober]]?

In (6) (Lakoff's 'Type C'), the interpretation is that the state of affairs described in the second conjunct is a **result** of the state of affairs described in the first conjunct, and extraction is again from the first conjunct only.

(6) What do people [[eat __ here] and [then get sick]]?

Lakoff develops an analysis of these correlations with the following two main components:

- Conjunctions such as English *and* are interpreted as expressing a variety of discourse relations between propositions.
- Patterns of extraction from coordinate structures are conditioned by the discourse relation expressed by the conjunction.

This is a radically different approach to Ross's, because the acceptable patterns of movement are described not in syntactic terms like 'coordinate structure' or 'conjunct', but in terms of descriptions of interpretations like 'result' or 'narration'. However, Lakoff only actually describes correlations between certain interpretations and certain patterns of extraction. These correlations (even if empirically accurate) are only a step toward a full theory, because we don't know why the correlations should be this way, or why they should be visible in this empirical domain.

We might hope that established theories of syntax and of discourse structure could help answer these 'why' questions. However, we then encounter a further property of this empirical domain:

- Different syntactic theories have very different analyses of coordinate structures, and make significantly different predictions about extraction from coordinate structures. No theory currently makes fully accurate predictions for the attested data.
- Different theories of discourse structure make significantly different predictions about the interpretation of conjunctions like *and*. No theory currently makes fully accurate predictions for the attested data.

We will see in Chapter 2 that there is an almost embarrassingly wide range of analyses of the syntax of coordination under active current development. Chapter 3 shows that most theories don't predict the noted diversity of extraction patterns, and that we have a choice between analyses which are too restrictive, or which are too liberal.

Formal theories of discourse are newer, and the range of analyses is narrower as a consequence, but there are still several approaches under current development. The labels we used above ('parallel', 'contrast', 'narration', 'violated expectation', 'result') are fairly widespread, and, in fact, have been defined in several slightly different ways (see

Chapter 5), but only one work (Kehler 2002) has attempted to relate reasonably explicit definitions of those labels directly to patterns of extraction. Kehler (2002), therefore, is the most direct antecedent of the discourse research reported in this monograph, but we will diverge from Kehler in several important respects, particularly in Chapters 5 and 6.

Because neither the syntactic theory nor the discourse theory has stabilized, there are many moving parts to consider when attempting to build a unified account of syntax and discourse in this empirical domain. This problem has sometimes been underestimated on both sides: syntacticians have often been content with a demonstration that extraction from coordinate structures is semantically conditioned, without attempting to show *how* semantics influences a syntactic phenomenon like nonlocal dependencies (e.g. Munn 1993, Johannessen 1998). And the question of extraction from coordinate structures barely registers in many works on discourse semantics (with the notable exception of Kehler 2002), because, well, that's syntax.

Perhaps as a result of this disconnect, there have been two main approaches to extraction from coordinate structures in recent work. One, which we will refer to as 'syntax calls the shots', aims to reduce patterns of extraction from coordinate structures to principled statements about constraints on unbounded dependencies in syntax. The other, which we will call 'discourse calls the shots', aims to explain restrictions on extraction from coordinate structures in terms of asyntactic statements about the interpretation of unbounded dependency constructions in specific discourse contexts. We will survey research to date in these two approaches in Chapters 4 and 5 respectively. The following is a very brief summary of our findings.

- Certain variants of the syntax-calls-the-shots approach can offer real insight into data related to (3)–(6), but this approach offers no real hope for addressing the heart of Lakoff's challenge, namely the relationship between extraction patterns and interpretation of coordinate structures.
- Claims, such as Lakoff's or Kehler's, that extraction patterns correlate with discourse relations such as 'narration' or 'violated expectation' are approximately correct, but no more than that. Accordingly, variants of the discourse-calls-the-shots approach such as Lakoff's (which build directly on such relations) cannot address the relationship between extraction and interpretation in detail, despite their initial promise. These approaches also cannot capture certain patterns which admit explanations in a syntax-calls-the-shots vein.

The summary just given is largely negative, but all is not lost. There are two ways forward, which jointly offer potential for preserving the strengths of the two approaches while avoiding the weaknesses. The first is simple: divide and conquer. The syntax-calls-the-shots approach can explain some patterns quite naturally, in ways which are beyond the scope of the discourse-calls-the-shots approach. So we should let syntax take care of those patterns. To give one example, there is some crosslinguistic variation with respect to the extraction patterns in (3)–(6). The most important variation concerns the availability of extraction from noninitial conjuncts that are interpreted as being part of a narrative, like (4). For instance, English allows this pattern, but French and German do not. A syntactic explanation of this difference will be sketched in Chapter 4, but it is unclear how this crosslinguistic difference could be pinned on an invariant fact about discourse structure, such as the

availability of narration-like readings for VP coordination structures, which is common to all three languages.

The second way forward comes from examining the theoretical status of discourse relations. Lakoff treated the different interpretations of coordinate structures as unanalyzed primitives, but there have been attempts since then to look for a principled reason why these particular relations feature in the way that they do. Kehler (2002) attempted this, but we will conclude in Chapter 5 that his approach is incomplete. We develop an approach in Chapter 6 that we believe to be more promising. This approach takes SDRT's analysis of discourse relations (Asher 1993, Lascarides and Asher 1993, Asher 1999, Asher and Lascarides 2003, Asher and Vieu 2005, *inter alia*) and coordination (Txurruka 2003) to generate hypotheses about common properties of interpretations of coordinate structures, and information-structural differences between these interpretations. These information-structural differences motivate the different extraction patterns that Lakoff noticed, as well as many nuances that are beyond his, or Kehler's, approach.

In sum, the questions that motivate this monograph have been investigated from a number of angles, for over fifty years, but to our knowledge, no previous work offers a truly integrative perspective on these questions, covering syntax, semantics, discourse, and the relations between these domains. We hope to demonstrate that progress in this area requires this integrative perspective. Our attempt to synthesize work across these domains has made us confront critical opportunities for comparison among established current syntactic and semantic theories. It has also substantially sharpened previous research questions, and implied a novel research agenda.

1.3 Structure of the volume

The first two chapters lay out the empirical focus of the monograph more precisely. **Chapter 2** begins by searching for definitions of 'coordination' and 'coordinate structure'. Section 2.2 proposes a first pass, roughly following Chomsky (1957), based on the propositional connectives \wedge and \vee , with subsentential coordination being derived from sentential coordination by ellipsis. However, Section 2.3 reviews evidence that not all coordinate structures can be reduced to sentential, or propositional, coordination, for instance because 'group-forming' NP coordination does not have a propositional analog (*Mary and John hugged* is not derived from *#Mary hugged and John hugged*). This suggests that coordinate structures are not semantically uniform: some are interpreted in ways which can be paraphrased as propositional conjunction or disjunction, while others are interpreted as groups.

In response to this, Section 2.4 discusses a range of analyses of the morphosyntax of coordination, developed on grounds independent of extraction. However, this section fails to identify a clear morphosyntactic 'hallmark' of coordination: no necessary and sufficient properties for distinguishing coordinate structures are identified. This conclusion chimes with the semantic map proposed by Haspelmath (2004), in which a range of different coordinate structures may be distinguished within a language, and certain types of coordinate structure may shade into adverbial and compounding structures.

Chapter 3 introduces data on patterns of extraction from coordinate structures. It is sometimes suggested that the patterns, and particularly the Coordinate Structure Constraint introduced in Section 1.2, is just the kind of hallmark of coordinate structures that

Chapter 2 failed to identify. Section 3.2 gives a fuller introduction to the CSC, and Section 3.3 describes many documented classes of counterexamples, which problematize the notion that the CSC is a hallmark of coordinate structures. Section 3.4 aims to dispel the possible impression that the patterns described in this chapter relate specifically to *and* and similar lexical items in Western European languages. While a proper typological survey of the CSC and its exceptions is not possible at this time (in part because descriptive grammars typically do not address such issues at all, and certainly not in the necessary depth), we can nevertheless show that similar patterns hold across multiple conjunctions in several genetically and areally unrelated languages. Section 3.5 discusses the scope of the CSC and its counterexamples across A', A, and head movement. Finally, Section 3.6 introduces a particularly complex case, the **SLF** construction found in many languages, but most widely discussed with reference to German and Dutch. This construction involves the interaction of syntactic and interpretive phenomena, which are informative about the syntax and discourse semantics of coordinate structures. Although there is a large literature on this construction, it is often ignored in discussions of Lakoff-style effects, so we devote extra space to it here.

The survey in Chapter 3 suggests that almost all cases of asymmetric extraction are specifically from VP coordination. This entails that, in order to make sense of Lakoff's data and the challenge that the data pose, we need to strike a proper balance between the general validity of the CSC and the range of apparent counterexamples concentrated in the special case of VP conjunction. In turn, this raises a general question, discussed in Section 1.2 and articulated more fully in Section 3.8: what kind of explanation should we seek for these patterns? The CSC, as formulated by Ross, clearly contains more than a grain of truth, but is strictly falsified by the patterns discussed in the second half of Chapter 3. On the other hand, discourse-based accounts such as Lakoff's suffer from several shortcomings. The question, then, is: who (if anyone) calls the shots?

Chapters 4 and 5 discuss this question from the perspective of syntax and discourse, respectively. **Chapter 4** begins with a discussion of the desiderata for a syntactic account in Section 4.2, and a brief introduction to mainstream locality theory, the theoretical context for the CSC, in Section 4.3. The main point, articulated in Section 4.4, is that mainstream locality theory has 'moved away' from the CSC since Ross's thesis. The processes of unification and generalization of locality theory in Chomsky (1973, 1981) implied increasingly clearly that the CSC was different. In response to this, Section 4.5 discusses locality theories which aim to incorporate the CSC and derive it from more general principles.

This establishes a choice point. Chomskyan locality theory essentially ignores the CSC, but the theories discussed in Section 4.5 integrate the CSC so tightly with the foundations of locality theory that there is only limited scope for addressing counterexamples. Hence, we need a third way. Section 4.6 discusses one possibility in detail, namely the possibility of relating patterns of extraction from coordinate structures to patterns of extraction from adjunction structures. The leading idea in this approach is to capitalize on similarities between ATB extraction and **parasitic gaps**, the other major class of cases in which a single moved element corresponds to multiple gaps. The approach has had a rocky history, mainly because Postal (1993) gave a thorough account of the many *dissimilarities* between ATB extraction and parasitic gaps. Our estimation is that Hornstein and Nunes (2002) provides the outline of a viable response to Postal's argument: parasitic gaps are more restricted than ATB extraction, and these additional restrictions are due to an additional syntactic factor

superimposed on an essentially identical movement configuration. The rest of Section 4.6 develops further predictions of this conjuncts-as-adjuncts analysis. The broad outline of these predictions (a slight reformulation of ideas in Postal 1998) is that extraction from initial conjuncts should be syntactically unrestricted, while extraction from non-initial conjuncts in VP conjunction structures should be as restricted, crosslinguistically and intralinguistically, as extraction from VP adjuncts. We show in this section that those predictions have some validity, but leave a string of unanswered questions.

The upshot of Chapter 4 is that certain options within syntactic theory can predict patterns of extraction from coordinate structures with greater subtlety than Ross could, without resorting to treating extraction from coordinate structures as *sui generis*. However, this does not, in itself, constitute a response to Lakoff's challenge. In fact, the patterns described in Section 4.6 are strictly orthogonal to those described by Lakoff. That is, it is a syntactic fact that extraction from initial conjuncts is often easier than extraction from noninitial conjuncts, and that noninitial conjuncts (even when they allow extraction) behave like weak islands. On the other hand, if Lakoff (1986) is correct, it is a nonsyntactic fact that 'narration'-like coordinate structures allow extraction from noninitial conjuncts, while 'violated expectation' and 'result' allow only extraction from initial conjuncts. These two types of observation appear to have irreducibly different statuses, but they both pertain to the same set of empirical phenomena.

Chapter 5 investigates Lakoff's analysis on its own terms. The heart of the analysis is a series of correlations between discourse relations and patterns of extraction. In Section 5.2, we outline Kehler's (2002) taxonomy of discourse relations. This taxonomy is based on seminal work by Hobbs (1979, 1985, 1990), who extends David Hume's philosophical ideas about the association of ideas to natural language discourse. We also provide a brief glimpse of how the Hobbs/Kehler program—originally intended for AI research—has led to fruitful research in semantics and pragmatics. Subsequently, in Section 5.3, we consider Kehler's (2002) formal definitions of some well-studied discourse relations that are relevant to Lakoff's analysis, namely OCCASION, RESULT, VIOLATED EXPECTATION, PARALLEL, and BACKGROUND. We discuss the challenges and shortcomings of Kehler's (2002) definitions and provide a glimpse of steps that have been taken to address some of these challenges.

Against this background, Section 5.4 evaluates the fit between these independently defined discourse relations and patterns of extraction from coordinate structures. Our conclusion is that there is no real scope for analyzing discourse relations as direct causal factors explaining the different extraction patterns that Lakoff identifies. There are two reasons. The first is that Lakoff's correlations don't give any immediate way to explain why VP conjunction allows this range of interpretations, but not others. The second, and more challenging, is that once we have explicit and reasonable definitions of NARRATION and RESULT, it becomes apparent that these relations do not stand in opposition to each other. Rather, they stand in an entailment relation: RESULT is a special case of NARRATION. This drastically limits the scope for explaining Lakoff's patterns in these terms: if RESULT is a special case of NARRATION, how are we to explain the apparent fact that RESULT doesn't allow the same extraction patterns as NARRATION? We conclude that discourse relations in their own right could not explain Lakoff's correlations, even though those correlations were originally stated in terms of discourse relations.

In the light of this conclusion, **Chapter 6** develops a hypothesis mentioned briefly by Kehler (2002), namely that the distribution of **topics** mediates the relationship between

discourse relations and extraction patterns, because only topics can be extracted from coordinate structures. Section 6.2 introduces, fleshes out, and critiques Kehler's hypothesis. The most important criticism is that Kehler overstates the role of topics. In some cases, particularly the examples with a 'narration' interpretation and extraction from a noninitial conjunct, topicality seems to drive patterns of extraction, but in most other cases, examples can be found where an element which is clearly not a topic can nonetheless extract.

This leads us to introduce SDRT, as a formal, integrated theory of discourse structure with well-developed analyses of discourse relations, topicality, and coordination. SDRT gives us the resources to state hypotheses about why topicality is important with some discourse relations, but not others.

In Section 6.3, we spell out the graph-theoretic basics that underlie SDRT's analysis of discourse structure, focusing on a property of discourse relations that is especially important: the distinction between **subordination** and **coordination**. Section 6.4 then introduces Txurruka's (2003) groundbreaking hypothesis that *and* is only compatible with coordinating discourse relations.

Next, in Section 6.5, we consider how logical forms of discourses are constructed according to SDRT, focusing on how discourse relations factor into the construction of topics, and a broadly similar object called a **common theme**. The crucial point is that some discourse relations create discourse units that stand in a particular structural relation to a (possibly implicit) discourse topic, and others give a particular structural status to the common theme, while still others have no explicit discourse topic or common theme.

With all of these elements in place, Section 6.6 considers the prospects of working out an SDRT-based analysis of extraction from coordinate structures, and Section 6.7 extends this analysis to coordinate structures with more than two conjuncts. In particular, we propose that there are four patterns to consider, corresponding to four classes of relation just described:

1. Subordinating relations cannot be expressed by coordinate structures in the first place;
2. discourse relations which refer in their semantics to a common theme require ATB movement in the syntax;
3. all other relations also permit asymmetric extraction from initial conjuncts;
4. discourse relations which stand in a structural relation to a discourse topic allow extraction of topical elements from noninitial conjuncts.

Chapter 7 rounds off the survey with an evaluation of the strengths and weaknesses of the different approaches surveyed, focusing on a comparison of syntactic analyses from Chapter 4 and the SDRT-based approach developed in Chapter 6. We also give a summary of the choice points and their interrelationships, and a series of open questions, intended as a stimulus to further research. The new SDRT-based approach to extraction from coordinate structures developed in Chapter 6 has implications for everything from the syntax of coordination through to the relationship between information structure and extraction. We think that it is a new way of stitching together pieces of syntactic and discourse-semantic analysis to give a promising unified whole. However, several questions need to be answered in future research before it can be properly evaluated. These include the following:

- What are the implications of this analysis for the rest of syntactic locality theory?
- What can we learn about the nature of islands from this approach?
- What is the bridging hypothesis linking information structure and patterns of movement?
- What are the possible loci of crosslinguistic variation of extraction, on this analysis?

1.4 What this volume is, and isn't

Although the empirical scope of the survey is narrower than those of many other surveys in the series, we believe that this is necessary in view of the theoretical breadth that the topic requires.

As surveys go, this one is quite opinionated. We have endeavored to do justice to the range of current theories of relevant areas of syntax and discourse semantics, as well as to the history of these ideas, but we have also chosen to focus on theoretical choices that, in our opinion, hold real descriptive advantages. There is a limited amount of original research in the survey, in the service of these opinions.

There are many things that this book is not a survey of. It is not a survey of:

- Coordination. Van Oirsouw (1987) is an excellent survey of the early work on the syntax of coordination, not yet surpassed in its coverage despite the significant advances in understanding since 1987. More recent monographs with a survey aspect include Zhang (2009), while Progovac (1998a,b) briefly surveys the range of then-current proposed coordinate structures, and Haspelmath (2004, 2007) are particularly useful as chapter-length typological overviews.
- Locality. Recent surveys of locality effects and locality theories include Rizzi (2013) and Boeckx (2012), as well as Chapter 2 of Truswell (2011).
- Information structure. Among others, the survey in this series by Erteschik-Shir (2007) is a comprehensive introduction to information structure and its interface with syntax.
- The relationship between competence and performance. There is a large body of work on sentence processing and island constraints (see many chapters in Goodluck and Rochemont 1992), and Harris (2011) has demonstrated that processing studies can help us understand the syntax and semantics of extraction from coordinate structures.
- Discourse structure. Zeevat (2011) and Jasinskaja and Karagjosova (2020) are excellent surveys of work on discourse relations, with the latter also discussing their impact on discourse structure. For an overview of discourse structure within SDRT, see Asher and Vieu (2005). For a basic guide to SDRT, see, e.g., Lascarides and Asher (2007), Altshuler and Schlöder (2019).

We will truncate discussions in all of these areas in order to keep our eye on the prize. We will see that some elegant and parsimonious current syntactic theories leave empirical holes that may be fillable by a discourse-semantic analysis of coordination. Discourse theorists, for their part, are just trying to develop theories of discourse structure, without paying particular attention to what syntactic theory would like it to do. How close can we get to a complete account just by joining the dots between bodies of research that don't typically

interact enough? How well do the different theories articulate? Where are the gaps, and what are the prospects for filling those gaps? These are the questions that we *do* aim to cover.

Finally, there are several other putatively syntactic phenomena where similar questions arise, including gapping and other ellipsis phenomena, and other 'island' effects. Any of these in principle could have been chosen as the empirical domain for this survey. We hope that the approach we take here to joining the dots can inspire similar work in those other areas.