COPiL CAMBRIDGE OCCASIONAL PAPERS *in* LINGUISTICS
Volume **13**, Article 3: pp. 63–97, 2021 | ISSN: 2050-5949

Towards a Conceptual-Semantic Model of Cross-Sentential Anaphora *

Roberto B. Sileo
The University of Cambridge

Kasia M. Jaszczolt
The University of Cambridge

Abstract The search for an antecedent in cross-sentential pronominal anaphora has given rise to various attempted generalisations and rules, none of which adequately capture the complexity of the phenomenon. In this paper we argue that the current views on the topic are all partially correct but they err in their polarization: they attempt either to stay on the side of quite free, unconstrained pragmatic solutions or they go too far in semanticizing or even grammaticizing the regularities. Instead, we propose to capture both the regularities and the diversity of the phenomenon by approaching it on the level of conceptual structure, understood as a radically contextualist semantic representation that allows more scope for free pragmatic inference, at the same time retaining the rigour of a formal, truth-conditional representation. We also report the results of our small pilot study that (i) strongly suggests that a wide array of factors are responsible for context-free, context-neutral, as well as contextually-biased anaphora resolution and (ii) leads us to questioning the need for further empirical search for regularities on the level of linguistic structure. The relevant factors include, but are not limited to, the salience of the concept associated with the specific lexical items employed, interlocutors' social and cultural assumptions, particular (including imaginary in the case of accommodation) situation of discourse, the distance on the memory line, and world knowledge. We finish by presenting sample representations of salient interpretations of relevant utterance pairs in the framework of Default Semantics – a radically contextualist theory that advocates the compositionality of meaning on the level of conceptual structures. We conclude by proposing that the ongoing debates between grammar-based and pragmatics-based solutions are futile in that they are all partly correct, and that their findings can inform a solution like ours, where pragmatics-rich semantic representation allows to accommodate the regularities or preferences they have uncovered.

* Research for this paper was partly supported by the Cambridge Humanities Research Grant *Rethinking Being Gricean: New Challenges for Metapragmatics* https://www.mmll.cam.ac.uk/rethinking-being-gricean. We are grateful to the audience of the American Pragmatics Association Conference AMPRA 4, Albany, NY for their feedback, to Andrew Caines for his advice on the quantitative analysis, to Elaine Schmidt for help with finding participants through the Cambridge Linguistics Research Study System, and most of all, to Ernie Lepore and Matthew Stone, the authors of *Imagination and Convention*, for inspiring us to revisit this interesting phenomenon from the conceptual-semantic perspective.

1 Introduction

The search for an antecedent in cross-sentential pronominal anaphora has given rise to various attempted generalisations and rules, none of which adequately capture the complexity of the phenomenon. While there is no doubt that context can affect the choice of the suitable antecedent, the debate is on-going as to whether the English language contains generalizations, other than specific concept- or specific function-based,[1] that can lead to a semantics with normative rules for cross-sentential reference assignment or, rather, the phenomenon is to be left to non-formalizable pragmatic heuristics. While computational linguists tend to lean towards uncovering language-internal regularities, Gricean pragmaticists prefer to pin any ensuing correlations and preferences onto heuristics that are said to govern communication, cognition and, generally, rational human behaviour. In this position paper we argue that the current views on the topic are all partially correct but they err in their polarization: they attempt either to stay on the side of quite free, unconstrained pragmatic solutions or they go too far in semanticizing or even grammaticizing the regularities. Instead, we propose to capture both the regularities and the diversity of the phenomenon by approaching it on the level of conceptual structure, understood as a radically contextualist semantic representation that allows more scope for free pragmatic inference, at the same time retaining the rigour of a formal, truth-conditional representation. The relevant proposition is what we call the *functional proposition* that pertains to this conceptual structure and as such combines information conveyed through different linguistics and non-linguistic means. We also report the results of our small pilot study that (i) strongly suggests that a wide array of factors are responsible for context-free, context-neutral, as well as contextually-biased anaphora resolution and (ii) leads us to questioning the need for further empirical search for regularities on the level of linguistic structure. The relevant factors include, but are not limited to, the salience of the concept associated with the specific lexical items employed, interlocutors' social and cultural assumptions, particular (including imaginary in the case of accommodation) situation of discourse, the distance on the memory line, and world knowledge. We finish by presenting sample representations of salient interpretations of the relevant utterance pairs in the framework of Default Semantics – a radically contextualist theory that advocates the compositionality of meaning on the level of conceptual structures. We conclude by proposing that the ongoing debates between grammar-based and pragmatics-based solutions are futile in that they are all partly correct, and that their findings can inform a solution like ours, where pragmatics-rich semantic representation allows to accommodate the regularities or preferences they have uncovered.

[1] For instance, it was found that in story-continuation tasks in which perfective forms of verbs of transfer were tested, such as in 'A handed a book to B', participants selected antecedent goals in indirect objects (in this example, B) significantly more often than sources in subjects (in this example, A). See Kehler, Kertz, Rohde & Elman (2008); Ferretti, Rohde, Kehler & Crutchley (2009). This is not the level of generalization we are interested in here.

The structure of the paper is as follows. In section 2, we present some highlights from the cross-sentential pronominal anaphora discussion, attending to the problems with the tendency to polarise the views as grammar/semantics-oriented vs. pragmatics-oriented. We start with a brief assessment of the solutions in coherence theory tradition, move to attention-based Centering Theory, and then assess more recent hybrid solutions. Next, we focus on dynamic semantics approaches that lead us to our own proposal. In the case of each selected approach, we pinpoint the main aspects that we hypothesise ought to be inherited in our own conceptual-semantic approach. In section 3, we present our small empirical source of ideas in the form of five case studies of cross-sentential anaphora, tracked through different context conditions, also addressing the methodological question concerning the utility and place of experimental methods in the search for relevant principles and heuristics that would cater for cross-sentential anaphora *en masse*. We conclude the section by proposing a tentative typology of factors that affect anaphora resolution that suggest the feasibility of a normative account. Section 4 offers a proposal of semantic *qua* conceptual representations in the framework of Default Semantics (henceforth DS, Jaszczolt e.g. 2005, 2010) that capture the diversity of factors on which anaphora resolution relies. Section 5 concludes with further thoughts on the feasibility of a model of cross-sentential anaphora resolution, the normative status of such a theory, as well as metasemantic remarks on the differences and convergences between computational and Gricean accounts.

2 Cross-Sentential Reference Assignment: Semantics and pragmatics

2.1 Linguistic rules or general heuristics?

The question as to what factors influence cross-sentential anaphoric reference assignment in English is a subject of ongoing debates. Who is expecting the quick confirmation in (1)? Does the pronoun 'he' refer to the president or to Jones?

(1) The president nominated Jones. He expected a quick confirmation.
(from Lepore & Stone 2015: 91)

In an attempt to show that utterance interpretation is linguistic-convention-driven rather than pragmatics-driven as originally envisaged by Grice (1989), Lepore & Stone (2015) suggest, following the proponents of attention theory and the so-called *grammatical role hierarchy* (see e.g. Grosz, Joshi & Weinstein 1995), that the language system itself contains a specific rule for cross-sentential reference assignment. In their view, in a context-free situation where anaphoric reference to the subject and to the object are both plausible, the referent of a pronoun in the subject position of an English sentence is preferably resolved by the immediately preceding subject. They argue, in this respect, that in the absence of contextual assumptions to the contrary, 'he' in (1) refers to the president. On the other hand, 'he' in (2) refers to Jones.

(2) Jones was nominated by the President. He expected a quick confirmation. (from Lepore & Stone 2015: 91)

In discussing examples (1) and (2), Lepore and Stone invoke the grammar-based principle of attention whereby the subject position is claimed to afford greater prominence than the object position as far as the candidates for the antecedent are concerned. Following attention-based theories, they regard syntactic structure as a good guide for salience ranking. They claim that, by convention, and in the absence of contextual biases, 'there is a preference to resolve a subject pronoun in one sentence to the subject of the preceding sentence' (Lepore & Stone 2015: 91). The principle of attention is explicitly syntactic: as such, it allows Lepore and Stone to identify the bias as a grammatical convention and combine coherence considerations with the explanatory role of the centre of attention, making them both conform to 'linguistic rules' that can be captured by the logical form. Building a larger picture of bound as well as demonstrative uses of pronouns, they argue elsewhere (Stojnić, Stone & Lepore 2017, 2020) that semantic values of pronouns follow the linguistic rules to which the context of use conforms. In other words, attention and coherence are governed by linguistic rules (see e.g. Stojnić et al. 2017: 532).

We begin by placing Lepore and Stone's stipulated linguistic convention in the context of the views that search for what can count as their underlying causes. First, theories of discourse coherence (e.g. Hobbs 1979, 1990, Kehler 2002, Kehler et al. 2008, Asher & Lascarides 2003) explain pronoun interpretation by invoking semantic as well as pragmatic factors such as world knowledge and context-driven inference. On the semantics-pragmatics scale of proposed solutions, they are often placed towards the pragmatics end in that they focus on the interlocutors' expectations of coherence where the latter is systematised as coherence relations such as, for example, Parallel or Result. Kehler et al. (2008) derive explanations from grammatical-role parallelism, subjecthood, or thematic roles that form such coherence relations. There are various heuristics that lead to interpretative preferences but since the heuristics are often incompatible with each other, common-sense, world knowledge and context have to lend a helping hand. These preferences jointly result in discourse coherence. Various typologies of coherence relations (Hobbs 1990, Kehler 2002, Asher & Lascarides 2003) arguably delve deeper into the reasons for coreference biases. So, coherence theorists allow a modest place to the syntax-based preference foregrounded by Lepore and Stone, focusing on the fact that there are other types of preferences that may take precedence.[2]

Next, Centering Theory (Grosz et al. 1995, Grosz & Sidner 1998, Walker, Josh & Prince 1998) offers an explanation from the attention state. It proposes a model of a relationship between attention state, understood as availability of referents, that corresponds to the degree of activation of discourse referents at a given point in discourse, the form of a referring expression, and the coherence as assessed for a relevant fragment of discourse (see Gundel 1998, Gundel, Hedberg & Zacharski

[2] See here Kehler's (2002: 143) discussion of Kameyama's empirical evidence concerning the prominence of the subject vs. the object position; Kehler et al. (2008) and Kehler & Rohde (2013), and the discussion below on the convergences in the discussed approaches.

1993: for an example of such a model). Through capturing attention states associated with referents in the processing of discourse (degrees of topichood), the theory accounts quite well for some regularities in discourse anaphora.

Coherence-based and attention-based views are very different but not incompatible. Focusing on inferential processes that establish coherence on the one hand, and focusing on topic relations and grammatical roles of referents on the other, offer very different explanations of discourse relations. But they offer probabilistic explanations that are arguably both needed: as Kehler & Rohde (2013) demonstrate, (i) expectations concerning coherence relations that result in a selection of the antecedent and (ii) the theories of topichood are both needed to explain the phenomenon.

We want to go further. Instead of fuelling the theoretical debates concerning the predictive power and mutual compatibility of the extant orientations, we offer a new solution in the form of the conceptual-semantic representation to which a wide range of factors contribute, but this contribution is structured and harnessed by associating the factors with the sources of information and types of processes of interpretation we identify, following the theory of DS. But before we get there, we should look at some hybrid solutions.

First, conventional associations should not be discarded outright. The question is whether these conventions are best considered to be linguistic conventions and as such represented in the logical form *à la* Lepore & Stone (2015) and Stojnić et al. (2017). Even though pronouns pick out the referent that is in the centre of attention and as such governed by some or other model of activation, and even though, on the other hand, coherence theory relies on an array of fairly reliable coherence relations, it seems that we have to 'kick the analysis one level up', so to speak, if we aim to capture how the interaction of all these factors results in the actual multidimensional heuristics applied by the conversational interactants – and, at the same time, if we aim to offer generalizations for what would otherwise be seen as situation-induced exceptions. While we agree with Stojnić, Stone and Lepore that the uncovered generalizations make the solution semantic *qua* rule-based, we do not think that it belongs with the grammar of English on that level of generalization.[3] We don't think that the solution is that simple. As they say, there are many different, often competing, mechanisms that contribute to pronoun resolution but for them 'their contribution is governed by rules, not pragmatic reasoning, and as such, should be formally represented in logical form' (Stojnić et al. 2017: 540) For us, being rule-governed does not guarantee that we are dealing with linguistic information, even in their generous sense of the term. Neither does semantic contribution equal contribution to the logical form, where the latter is reserved for the *linguistic semantic representation*, that is the output of processing of the syntactic structures of natural language *tout court*. On the other hand, the fact that they are looking for a more precise, formalizable solution than, say, neo-Gricean heuristics, and as such a conceptual-semantic solution, is what puts us on broadly the same

[3] See footnote 1.

path: conventions have been grossly understated in Gricean pragmatics and it is time to put this right.[4]

While there is an undeniable factor to do with the prominence of the subject position, departures from this pattern also follow rules with pretty strong predictive power. So, it is often pointed out by coherence theorists that *semantic focus*, that is, 'comment', 'new information' in the topic/comment distinction, is more often associated with the position of the object than that of the subject. But if this is so, then objects make good candidates for antecedents. On the other hand, Centering Theory stresses the fact that the subject is more often a *psychological focus* (as it is the so-called 'backward-looking focus' in discourse) and as such is a good candidate for an antecedent. As Gundel (1999) rightly points out, both observations are partly correct. The widely quoted example in (3a,b) (from Winograd 1972: 33) illustrates the case, whereby the referent of 'they' is normally easily resolved as the city council in (3a) and as the demonstrators in (3b).

(3) a. The city councilmen refused the demonstrators a permit because they feared violence.

b. The city councilmen denied the demonstrators a permit because they advocated revolution.

Be that as it may, inference based on commonly shared knowledge of the world captures the reasons why we attribute the fear of violence to the city officials and the plans of violent action to the demonstrators. And this common inference can easily be made into a heuristic as a coherence relation of Explanation. So, structure-driven prominence is not immune to other generalizable factors and is often regarded as a by-product of a more general heuristic. The task is to show how these heuristics can lead to a semantic representation without compromising the predictive power accumulated by different contributors on all sides of the debate – and in the spirit of the growing consensus that they are all partly right.

Lepore & Stone (2015: 88), in their homage to linguistic conventions, claim that in the case of anaphora that can, without violating the rules of grammar, be resolved in more than one way, the explanation ought to proceed along the lines of a disambiguation, or 'finding the right reading of the utterance, understood as a grammatically specified pairing of form and meaning'. By doing so, they take a stance against Gricean explanations of discourse that are grounded in the recognition of speaker intentions. While the metasemantic question concerning the compatibility of convention-based and intention-based accounts is pursued by one of the authors elsewhere (Jaszczolt 2016b, 2019), what interests us here is the question of the adequacy and relative power of such grammar-based and pragmatics-based explanations of the phenomenon. To repeat, even if indeed there were, hypothetically, empirical support for the greater prominence of the subject as the antecedent in examples sharing the structure with (1), there is a long way from there to demonstrating that reference assignment conventions that would apply across the board

[4] On neo-Gricean heuristics see e.g. Levinson (2000); see also Jaszczolt (2019) for a discussion of what counts as 'being Gricean' on the brink of the 2020s.

rather than to specific types of constructions are to be associated with the *power of grammar* – grammar that produces all the possible readings, leaving pragmatics to disambiguate. To repeat, while preferences to do with types of constructions have been attested and are, indeed, likely to occur in virtue of the semantics, pragmatics, or standard discourse-function of the expression type, there is no reason to hypothesise a grammar-based theory of cross-sentential anaphora resolution. But first, we will open this prominence to theoretical and, through our case studies, also some empirical scrutiny.

A disclaimer is in order at this point in that, unlike computational approaches to which the discussed views mostly belong, we will not be offering probabilistic models; neither will we search for experimental support. These have been developed extensively in the past four decades or so and have not led to conclusive results. Instead, we search for a semantic – compositional, truth-conditional, but at the same time cognitive rather than linguistic-semantic – representation that results from the heuristics applied by speakers on particular occasions. Unlike computational semantics, we do not propose to offer a ranked list of possible interpretations that are of use in training artificial intelligence systems. We search for a model that will give us predictive power for cross-sentential anaphora resolution as performed by human agents who use assumptions about speaker's intentions, founded on, in ordinary contexts, the Gricean principle of cooperative interaction. This anaphoric link will pertain not to the uttered sentences but, following the contextualist assumptions, to the primary message the interlocutors associate on a given occasion with the chunk of discourse. In this sense, we are Gricean – but, as will become obvious in section 3, only in this general sense that defines the object of analysis as the total impact of an utterance in a given situation ($\text{meaning}_{\text{NN}}$). Whether the uncovered generalizations can be easily applied to machine learning is a separate question that will not concern us here, although, bearing in mind an example of a successful implementation of neo-Gricean heuristics in Optimality Theory Pragmatics (Blutner & Zeevat 2004), we remain optimistic about the possibility of such an application. Since, arguably, this difference of objectives does not impact adversely on the search for principles for cross-sentential anaphora resolution, in that even context-free pairs of sentences that are of much greater interest to computational linguists as a source of rules are *de facto* assessed in a made-up, imaginary, or default context (on the proviso that one follows the contextualist assumption about meaning and truth conditions, which we do), we can proceed without committing a fallacy of talking at cross-purposes.

In brief, the grammar/semantics/pragmatics debate concerning the provenance of the bias in cross-sentential pronominal anaphora resolution is partly rooted in some differences of assumptions and objectives that we briefly mention in our earlier disclaimer. While computational linguists aim at uncovering ambiguities and putting the disambiguating information into linguistic rules of some kind, be it grammatical or semantic, discourse-pragmaticists adopt a larger dose of psychologism[5] and aim at uncovering the path that leads from speaker intentions to speaker

[5] On psychologism in pragmatics see Travis (2006) and Jaszczolt (2008).

meaning as recovered by the addressee (or, on more recent approaches, meaning jointly constructed by them, see Arundale 1999, 2010). But we can borrow from both in a 'positively eclectic' fashion. We can broadly follow the tradition from Winograd onwards in the search for the semantics that is sensitive to such options of reference resolution, but search it with the ultimate goal of finding out more about human cognition and its structure, pursuing the semantics of conceptual structures. Having uncovered these, we can move 'down' again to the level of grammar and semantics of natural language and explain cross-sentential anaphora through inheritance in that the thought processes captured in conceptual structure will also pertain to their vehicles to which natural language belongs. We return to these questions in section 3 where grammar and information structure are regarded as separate sources of information that contribute to a common conceptual representation.

2.2 Semanticizing anaphoric preferences: From DRT and SDRT to 'factors galore'

Lepore & Stone (2015: 2) justify inflating linguistic conventions and grammar by their assumptions that (i) '[p]ragmatic principles are fundamentally different from the principles of semantics' and, relatedly, that (ii) linguistic conventions that govern the meaning of expressions in natural language ought to capture what used to be delegated to conventions about agents' practices and rules (or: 'essential conventions', Strawson 1964). But, as we have pointed out, this assumption is not uncontroversial. Grammaticization of speaker's referential choices is a radical way of ascribing them to semantics – there are others. An alternative take on the question, situated somewhere between the Gricean intention-driven and Lepore and Stone's convention- (and grammar-) driven explanations, would be to afford semantic representations the status of mental representations *à la* Discourse Representation Theory (DRT, Kamp & Reyle 1993), allowing pragmatics some input that can, in principle, draw on conventions or on intentions. That would mean starting with the assumption that semantic theory ought to produce representations that have cognitive reality; semantic theory is not in the business of showcasing the power of the language system but rather the power of (human or other) agents in creating successful communication. The cluster of accounts we discuss next attempt precisely that. We discuss them because they provide groundwork to our own, more 'pragmaticised' solution.

Cross-sentential pronominal anaphora is a phenomenon that, in many ways, provided a *raison d'etre* for dynamic approaches to meaning, or at least a trigger for their development.[6] Kamp and Reyle start by pointing out that pronouns can be genuinely ambiguous between two or more resolutions, or they can even be uninformative whereby the interpreter is not able to arrive at a set of plausible alternatives. We will be interested here in the sources of information that help in the assignment of the suitable referent. Kamp & Reyle (1993: 67) observe that

[6] On philosophical foundations of dynamic semantics see Dekker (2012).

> [O]ften when the interpreter is able to reach a decision as to what the intended antecedent is, he does so on the strength of what he knows about the kind of situation that is being described and not exclusively on the basis of his knowledge of grammar or linguistic meaning. Such extraneous clues, which involve world knowledge rather than knowledge about language, are extremely difficult to analyse in any detail, and would require a much more elaborate theoretical framework than we could develop here.

We will adopt two aspects of this observation in our further discussion: first, that the intuition of a genuine referential ambiguity in neutral or otherwise unhelpful contexts is a fact of conversation, and, second, that, where a unique interpretation is recoverable in discourse, a satisfactory model of utterance interpretation that aims at predicting it correctly ought to incorporate an account of non-linguistic information as a component that can contribute to the derivation. The first aspect will be useful in the pragmatics vs. grammar argument and the latter in justifying a contextualist extension of DRT in the form of DS – both attended to in section 3.

Next, using the DRT framework, van der Sandt (1992, 2012), in his attempt to unify presupposition projection and anaphora resolution, advocates a 'preference order' of admissible antecedents that relies on the distance going up the projection line[7], but also on other factors that are far more difficult to formalise. This is an example of a 'soft' formalization of the grammatical preference order. Presuppositional expressions are regarded as anaphoric expressions and as such they are handled by the same resolution mechanism as anaphoric expressions – including discourse-anaphoric pronouns that we are interested in here. Presuppositions differ from standard anaphors in that they contain some semantic content and as such do not always require an overt antecedent; they can be accommodated.[8] This view of *presupposition projection as anaphora resolution* makes use of anaphoric binding as a mechanism that accounts also for what on other accounts, notably those of Stalnaker, Karttunen, Heim, Beaver i. a., is called *presupposition satisfaction*.[9] Presupposition projection, dubbed 'a repair strategy' (van der Sandt 1992: 345), allows for an anaphoric link to an antecedent that is made up, so to speak, on the basis of information inferred from context – hence we dubbed it here 'soft' formalization.

This account is useful for our argument in that it allows for a considerable interaction of preferences. First, it shows that the mechanism of going up the projection line can in some cases better account for the salience of antecedents than the uniform convention-based rule. For example, in (4), 'he' would normally refer to Harry's child even though the presuppositional inference disappears – and even though the nominal occupies the object position that, according to some accounts discussed earlier, comes with lower prominence.

[7] The 'projection line' is a sequence of discourse representation structures (DRSs) that form a path in the search for the accessible antecedents, going outwards from more to less embedded positions.

[8] On the concept of accommodation see Lewis (1979).

[9] See e.g. van der Sandt (2012) for a comparative introduction.

(4) Harry does not have a child. So he cannot be on holiday.
(from van der Sandt 1992: 335)

This preference can, of course, disappear in biased contexts but is otherwise quite robust. Van der Sandt explains it as follows. Anaphoric material is embedded in a DRS and then climbs up the projection line until the proper antecedent for binding can be found. According to this proposed mechanism, binding takes place at the nearest accessible position. In example (4), this position is occupied by the discourse referent for 'a child', present in a negated embedded DRS.

But projection alone will not help us with (1); in (1), both proper names either introduce discourse referents at the same level of representation, or, at the very least, have to emerge at the same level in the DRS by the time the second sentence is processed. This brings us back to the utility of this solution. Van der Sandt's ranking of the admissible interpretations according to a preference order testifies to the need for keeping the scope of relevant factors large and open. It makes use of various constraints such as full or partial matching,[10] relative distance along the projection line, discourse principles, and non-linguistic knowledge. For example, in (5), it is non-linguistic knowledge that leads to the salience of the binding to the object – again, in spite of the fact that there is no existential presupposition there.

(5) If John has a son, he will be at college now.

Now, accommodation is a well-acknowledged interpretive option where reference can be resolved through pragmatic bridging, as in (6), where 'she' quite unambiguously refers to the Queen.

(6) Britain benefits from being a monarchy. She brings the country good revenue from tourism.

Context, intonation, or information structure can make a supplied referent take precedence even over otherwise well-suited candidates as in (7) where 'she' is likely to refer to a female other than Meghan, most likely the Queen. Block capitals stand for emphasis.

(7) If Meghan and Harry remained abroad, SHE wouldn't like it.

It appears that if we want to capture every ensuing anaphoric link, we have to go quite 'pragmaticky' in our semantics.

So, although anaphora resolution through a projection line supplemented with the list of constraints are a step in the right direction, Kamp and Reyle's claim that including pragmatic constraints in a formal account is beyond the remits of the theory remains in place. Pragmatic factors are apt to 'interfere' (van der Sandt 1992: 362), or, as we would say in more positive terms, 'contribute', on an equal footing, to the interpretation. We therefore have some support here for a quest for

[10] An example of a partial match is 'If Anna has a new car, her car has a flat tyre.' The flat tyre can belong to her old car that she extravagantly abandoned or to the new one.

a comprehensive, contextualist account of anaphora resolution; most importantly, we have the tenet that pragmatic constraints are at least as important as the object they constrain. This will become important in section 3, where our case studies suggest that variety of types of influences on anaphora resolution are to be embraced rather than streamlined into ill-fitting generalizations. It will also be important in section 4 when we discuss the interaction of various interpretive processes in the recovery of the referent.

DRT has been further extended to incorporate rules of coherence. In their Segmented Discourse Representation Theory (SDRT), Asher & Lascarides (2003) offer a set of rhetorical structure rules that exemplify the overall presumption of discourse coherence and are said to guide the selection of the antecedent as part of the discourse's 'glue logic': the logic of information packaging (as separate from the logic of information content). They foreground the importance of discourse coherence for examples such as (8) and (9) where a pragmatically salient interpretation requires binding of 'he' to 'Bobby' in (8) and to 'Billy' in (9).

(8) Billy called Bobby a coward. He cried.

(9) Billy called Bobby a coward. He is a well-known bully.

The rule of maximisation of discourse coherence, supplanted with a series of rhetorical structure rules such as Result (in 8), Explanation (in 9), Narration, Elaboration or Background, building upon information from the lexicon,[11] explains the route to the preferred interpretation. The details of the account will not interest us here; what we need is the metatheoretic observation that conventions, in the form of rhetorical structure rules that spell out the non-monotonic glue logic, are effectively shortcuts through inference, are subject to preference order, but at the same time are subordinated to the pragmatic principle of the maximization of coherence that makes use of the logic of information content – or, in other words, of what is in the lexicon and in the structure.

Now, adding this set of rules to account for cross-sentential referential links may appear formal enough to count as only a small step removed from the 'linguistic conventions' solution whereby they are placed in the grammar itself. After all, by claiming that cross-sentential pronoun resolution derives from antecedents that remain at the center of attention in coherent linguistic interaction, Stojnić et al. (2017: 536) suggest that, 'as a matter of language', Narration makes antecedent subjects more prominent than objects, while Explanation makes antecedent objects more prominent than subjects. But there are, of course, cases where Narration does not necessarily raise subjects to prominence; they are present even in our limited case studies in section 3. And there are cases, as illustrated in (9), where Explanation does not necessarily raise objects. It is this multi-faceted interplay of the output of the logical form and the external factors that will interest us in pursuing a new approach that will give the diversity of sources of information what is due.

[11] See also Asher (2011) on rules and presupposition affecting the lexicon.

Needless to say, another way to handle (8) and (9) would be to employ default scenarios such as, say, children tend to cry when they are hurt and bullies tend to hurt people who are vulnerable. Note that the Result reading of (8) and the Explanation reading of (9) are only available after anaphora resolution; they are available following bridging in reasoning that necessitates precisely such default scenarios. To top it up, there is nothing to stop a person from uttering (8) in the context in which it was Billy who cried – for example because Bobby refused to join in playing a prank on a teacher. Crying with disappointment or anger, he retaliated by calling him a coward. *Mutatis mutandis* for (9). So, again, we have diversity of factors galore.

The question is, whether default scenarios will not suffice on their own: do we need the rules of coherence if they still leave the mappings quite messy? If our objective is to set out a list of factors that contribute to discourse anaphora resolution, then, arguably, any further abstraction over default scenarios falls outside our remit, its independent value for computational implementation notwithstanding. To wit, what we need to capture is the situations in which defaults do and do not apply, not the fact that rhetorical structure rules can generalise over selected cases to complete formal semantic representations; the latter is, to repeat, an intra-theoretic aim that is of more use in computational semantics where psychologism is not the main consideration. So, in preparation for setting out the positive proposal in section 4, we flag here the importance of default scenarios but not of any further abstractions over them. We take from SDRT three important lessons, namely that (i) conventions become apparent when the lexicon and the structure are interpreted; that (ii) they have to be treated seriously as contestants against Gricean intentions; and that (iii) they lead to further abstractions. But, to repeat, we can rest with default scenarios as these are more likely to lead us to a psychologically plausible contextualist representation of the variety of anaphoric links executed by discourse interactants.

In sum, accounts of cross-sentential anaphora offered in the dynamic-semantic tradition combine the methodological rigour of ascribing as much to the linguistic structure as descriptive syntax allows, with rigorous, step-by-step departures through ordered preferences on the cline of candidates for anaphoric binding and accommodation or, alternatively, through an appeal to general rules capturing cross-sentential coherence. To go further, one can search for a more precise 'preference order' for accommodation or for preferred rules of coherence, or, alternatively, instead of treating these regularities as 'one step removed from the grammar', concede that grammar is only one of the factors, all to be treated on an equal footing, and that the pool of candidates for antecedents includes those obtained through binding (with the distance as an important aspect), through accommodation (with coherence rules), socio-cultural conventions, degree of activation or topichood, situation-driven inference, among the important contributing factors. We conjecture that weighted sources of information pertaining to these regularities that contribute to the contextually preferred interpretation can give us a much needed model and, ultimately, possibly a normative account if the ordering itself turns out to obey some higher-order principles. In what follows, we offer first steps

towards such an account. But before we proceed, we report the results of our small pilot survey that helps us establish that such a pragmatics-rich, multi-factor analysis is indeed the best way to proceed and sheds some light on the different correlations that emerge. The specific format of case studies, whereby we take on board individual instances of cross-sentential anaphora but investigate them across suitably contrasting contexts and test them on a considerable number of consultants allows us to achieve just that.[12]

3 Of Subjects and Objects: Five case studies

Since we are offering here a position paper, its main purpose is to advocate the most adequate delimitation of the term 'semantic representation' to capture cross-sentential pronominal anaphora resolution, and, to some extent, also the mechanisms by which it is arrived at. We do so largely through theoretical argumentation and conceptual analysis. Position papers either precede experimental corroboration or render the latter superfluous in virtue of the logical progression of the argument. Our proposal sits on the fence but veers towards the latter. The results of our pilot study, in the form of seminal individual cases of anaphora resolution that we subjected to contextual modification and presented to our consultants, suggest that proliferation of surveys to test biases would be futile. We have believed from the start that instead of the quest for principles of coherence, rhetorical structure rules, activation-based preferences, grammar- or semantic-role based rules, to name only a few of the generalizations, all reasonably successful in themselves, it is more promising to identify the sources and processes that lead to the reference resolution and as such, attempt to construct an algorithm based on such interactive semantics of conceptual structures. The diversity of the relevant factors must not be shunned.

In section 2 we questioned the rationale for placing disambiguating reference assignment rules within the structure of English. Next, we conjectured that none of the extant theories of cross-sentential anaphora resolution go far enough into accounting for the myriad relevant factors and as such their predictive power will remain attenuated. In this section we explore this phenomenon from an empirical standpoint, following the behaviour of some individual cases of cross-sentential anaphora in and out of context. Our methodological assumption is that for the task at hand, namely shedding more light on the diversity of choices and factors, an in-depth analysis of individual cases is preferable to multiplying already extant quantitative surveys. Our results indicate that, in context-free settings, binding to the object position is much too common to propose a preference rule for binding to subject position. As we explain in what follows, we have also observed different degrees of entrenchment across the group of consultants of the context-free assignments when subjected to contextual modulation. The diversity of preferred choices, as well as the consultants' comments, allow us to propose a tentative typology of

[12] For an application of this method in the discussion of temporal ordering in discourse see Jaszczolt & Sileo (2021).

salient factors and, most importantly, proceed to the representation on the level of conceptual semantics that allows for a suitable lexicon/structure/pragmatics mix.

3.1 Questions, predictions, and hypothesis

The main rationale was to gain some insight into the relative salience of object and subject position as anaphors in neutral, context-free scenarios that in principle allowed for both the subject and the object to act as antecedents. We also wanted to test the entrenchment of these interpretations, across a reasonable number of consultants, when subjected to subject- or object-biasing contexts. This was expected to allow us to shed light on the claims discussed in section 2 concerning the rules and generalisations in cross-sentential pronominal anaphora resolution. We tested five individual cases of cross-sentential anaphora in three different linguistic contexts. In Questionnaire 1 (Q1), we looked into whether native British English speakers were biased towards either subject or object reference assignment in a context-free task that could, in principle, lead either to subject-based or to object-based resolution with no indication either way. To repeat, we tried to keep the scenarios neutral, although, as we explain in the discussion, potential reliance on default, imaginary scenarios means that the route to the context-free anaphora resolution may in itself rely on a variety of factors. In Questionnaires 2 and 3 (Q2, Q3), we investigated whether, and if so, how easily, consultants' responses can be skewed towards the subject- and object-based resolution. We hypothesised that if there were a subject- or object-biased reference assignment rule in the English language itself, consultants could be expected not only to be clearly led to either subject- or object-based resolutions in Q1 but also that these interpretations would be quite entrenched in biasing contexts in Q2 and Q3. Alternatively, if English did not contain either rule, consultants could be expected to reach either subject or object interpretations in Q1 and, once again, predominant subject and object interpretations in Q2 and Q3, respectively, due to rather unconstrained context-sensitivity. Following our argumentation in section 2, our hypothesis was that context-free cross-sentential reference assignment is not clearly subject- or object-oriented and we used our five case studies to pilot-test this assumption. The diversity of responses to each of the case studies was then deemed sufficient to proceed with the pragmatics-rich analysis of the phenomenon on the level of conceptual structure in that the more stringent posited correlations discussed in section 2 would have to relegate them to the category of counterexamples. In-depth individual case studies can sometimes reveal what large-scale surveys hide through calculating means. If the objective is to provide a psychologically plausible, normative account of how interlocutors resolve cross-sentential anaphoric links in discourse, the power of the theory has to be such as to account for our standard cases. Finally, a more detailed scrutiny of the effects of contextual bias applied to each case in Q2 and Q3 allowed us to propose a tentative typology of the sources of information for pronominal anaphora resolution, as well as some insight into their relative weight, utilised in section 4 for building conceptual-semantic representations.

3.2 Questionnaire design, participants, and procedure

As mentioned above, we tracked the behaviour of cross-sentential nominal anaphora in five case studies across three different context conditions presented in three separate questionnaires Q1, Q2 and Q3: context-free, contextually biased towards the subject and contextually-biased towards the object, respectively. The questionnaires themselves were designed for a broader task that elicited data on context-dependence of three different phenomena. The other two, not discussed in this paper, were temporal ordering in discourse and context-dependence of non-literal expressions. The phenomena were collectively tested so that they could function as mutual distractors. In addition, three further distractors were added, in the form of quantifying expressions that can give rise to scalar effects.[13] In order to license potential ambiguity, the subject and the object of the first sentence in each case study displayed consistent number and/or gender features and an effort was made to avoid the existence of strong default scenarios. The case studies are listed below as CS1-CS5. Notably, CS3 was Lepore & Stone's (2015) example discussed above as (1) in that we particularly wanted to check the validity of their intuitions.

CS1 The granddaughter congratulated the grandmother. She now started a new challenge.
CS2 The teachers failed the students. They felt disappointed.
CS3 The president nominated Jones. He expected a quick confirmation.
CS4 Kate brought Lilian to the meeting. She made a donation.
CS5 Tom invited Andrew. He enjoyed the party a lot.

In Q1 (context-free), the instruction was to describe in full detail the specific situation that the consultant first thought about when reading a short text, followed by an example. Q2 and Q3 (the contextualised versions) asked to describe the situation thought about in more detail. Next, the consultants were asked to state what they imagined might have happened next. On the one hand, detailed descriptions of the situations that consultants encountered were expected to reveal whether they resolved the anaphora via subjects or via objects. On the other hand, consultants' comments as to what might have happened next were requested in order to reveal subject or object interpretations in the event the description of the imagined scenario alone did not make the choice sufficiently clear. The pairs of sentences in each case study and the contexts provided in Q2 and Q3 are listed in Table 1 in the Appendix.[14]

Ninety native British English speakers recruited through University of Cambridge. Only non-linguists, who had been residents in the UK for at least three years, participated in the study (Q1: 13 female and 17 male; Q2: 21 female and 9

[13] We used the RAND function on Excel to randomise the presentation of our case studies and we manually manipulated the results (once) in order to avoid the consecutive presentation of two similar case studies.

[14] To test the efficacy of the questionnaire design, each draft questionnaire was completed by two volunteers who met our prescribed requirements. No material changes were made.

male; Q3: 21 female and 9 male).[15] In order to comply with the ethics rules and regulations of the University of Cambridge (and also to confirm eligibility), consultants were requested to sign a customary participation consent form. Consultants were paid £6 sterling as compensation for their participation. Electronic versions of Q1-Q3 were circulated which the consultants were asked to fill out in their own time and email back. The nature of this study (five individual case studies, each aimed at providing qualitative insight into the diversification of factors affecting the particular case of anaphora resolution) did not require testing the speed of responses.

3.3 Findings and preliminary discussion

Consultants' responses were classified as either subject-bound (S) or object-bound (O) interpretations. For example, in connection with CS1, consultants provided explicit S responses in (10) and (11) and explicit O responses in (12) and (13).

(10) The granddaughter asked the grandmother how to solve a problem. The grandmother found the solution and the granddaughter thanked her for her help. **The granddaughter moved onto the next task she had to complete.** (Q1/C15/CS1)[16]

(11) '...**The granddaughter** had set herself a series of fitness goals to complete, this being one of them, and having taken part in the run **she** looked ahead to her next goal ...' (Q1/C10/CS1)

(12) 'The grandmother had just completed a physical activity of some sort. **The grandmother now set herself a new goal to achieve.**' (Q1/C2/CS1)

(13) '**The grandmother** had been finding it difficult to get enough exercise...After a few months of effort she managed to complete her first park run without stopping. The granddaughter is a keen runner so is pleased at her grandmother's choice, and proud of her for managing it. With the granddaughter's advice, **she** then set herself a new target to run the distance in under an hour.' (Q1/C26/CS1)

[15] Consultants' ages ranged from 18 to 28 in Q1, from 18 to 23 in Q2, and from 18 to 35 in Q3. Consultants were recruited through the Cambridge Linguistics Research Study System, as well as through advertising at the University of Cambridge.

[16] We identify consultants' responses by the relevant questionnaire (Q), consultant (C), and case study number (CS1 - CS5).

Where consultants' responses were neither linguistically nor contextually explicit, we resorted to the strong linguistic and/or contextual clues that consultants provided in order to infer S or O interpretations. For instance, in connection with CS2, we considered the responses in (14) and (15) to represent, respectively, S and O interpretations. Relying on strongly implicated reference was deemed preferable to rejecting such answers as non-classifiable (N/C).

(14) 'A question comes up in an exam paper that the teachers had not prepared the students for, and as a result very few of them know how to answer the question properly and may fail the exam. **The teachers feel responsible for this**.' (Q1/C19/CS2)

(15) 'A number of lecturers did not show up on a certain morning, with no explanation or warning. **The student subject representative wrote an angry email on behalf of the students to the faculty**.' (Q1/C3/CS2)

Responses that made it impossible to determine which of the two interpretations had been obtained, as well as responses that referred to the two individuals (or set of individuals) in each case study, were regarded as N/C, as exemplified in (16) and (17) with reference to CS5.

(16) '**Two socially awkward guys** are at a friend's gathering… **Everyone** is playing games or pool. They go back to their rooms (they are housemates) slightly drunk and start supervision work for tomorrow.' (Q1/C4/CS5)

(17) 'Andrew was Tom's neighbour. They were friendly with each other but they did not share much in common. Tom considered Andrew boring but reluctantly invited him in the back of his mind hoping that he would be busy. **Both Andrew and Tom had a great evening**.' (Q1/C30/CS5)

As predicted, individual responses did not reveal consistent interpretative patterns. As we had anticipated, there is no evidence to support the view that cross-sentential binding preferences are dictated by the grammar – neither in Lepore and Stone showcase example (our CS3) nor in our remaining individual studies. In CS1, the overwhelming object bias in the context-free environment (86.7 per cent) remains almost stable when contextual object-biasing is introduced in Q3 (83.3 per cent, the difference of one response), which suggests that in Q1 either a salient, presumed context was supplied consciously, or a default scenario tacitly affected the interpretation. Subject-biasing in Q2 was, however, successful: it decreased the selection of binding to object by 63.4 per cent as compared with Q1. CS2 displayed a very different pattern: there was no clear subject or object bias in the context-free scenario (40 per cent vs. 56.7 per cent), and biasing contexts in Q2 and

Q3 were both successful in altering the preferences (subject-bias of 86.7 per cent in Q2 and object-bias also of 86.7 per cent in Q3). Lepore and Stone's showcase (CS3) displayed yet another pattern: it indeed displayed the subject bias in Q1 (albeit only 76.7 per cent), further increased to 93.3 per cent in Q2, and was not very successfully affected by an attempt at a biasing towards the object in Q3 (still 60 per cent of respondents selected binding to subject). But among out five case studies, this was the only one that displayed this pattern. CS4 and CS5 followed CS1 in an overwhelming object bias in Q1 (80 per cent and 73.3 per cent, respectively, although CS5 differed in producing a substantial number of N/C responses such as plural referents) but, again, differed somewhat in their sensitivity to attempts at contextual biasing: although both in CS4 and in CS5 the object-bias decreased not only in Q3 but also, curiously, in Q2, the intensity of contextual bias displayed different patterns (CS4 to 56.7 per cent and 63.3 per cent and CS5 to 60 per cent and 36.7 percent, respectively). Table 2 and Figure 1 in the summarise the results obtained for each case study across the three scenarios.

It goes without saying that we cannot take these mixed preferences to mean that English grammar itself favours out-of-context object-bound reference assignments. Rather, the fact that even our small case studies display preference for subject, preference for object, or pretty much evenly mixed preferences context-free, as well as different degrees of entrenchment when contextual biasing was attempted, supports an explanation in terms of a set of relevant factors that affect reference assignment – and as such contribute to the truth-conditional representation. Our five case studies had a quantitatively modest but at the same time qualitatively important task to perform, namely to exemplify that subject- and object-bound interpretations do not appear to be grammar-driven *tout court*; if the thematic role or syntactic position have a role to play, this role fits in the larger inventory of factors that all apply on an equal footing – or at least all apply with a certain weighting that, we hope, future large-scale corpus studies could help reveal. The first step attempted here is to highlight the importance of this messy state of affairs and a danger of unwarranted and hasty generalizations.

The next step is the pursuit of the factors that are relevant for discourse anaphora resolution, and here the theories discussed above give us some of them, and the case studies allow us to posit some more. Ultimately, this will allow us to pursue a contextualist, truth-conditional, semantic, but at the same time conceptual, representation of cross-sentential anaphora that defies the DR-theoretic claim cited above that the phenomenon itself lies beyond formalisation. To achieve this, one has to show how lexical items and sentence structure, in combination with other identified sources of information, trigger subject- or object-bound reference assignment – both in and out of context.[17] We will attempt this in section 4, following the search for the inventory of relevant factors in the remainder of section 3.

[17] 'To show how' is ambiguous between presenting the derivation and showing what components, and perhaps also through what processes, contribute to a compositional representation. In the DS-theoretic account pursued in section 4, we will opt for the latter. On extending DS to full derivation see Parikh (2016).

3.4 Further discussion: towards an inventory of factors influencing anaphora resolution

As expected, cross-sentential reference assignment appears to be sensitive to a variety of factors even in the absence of context, and the sensitivity to attempted contextual modification varies from case to case. The diversity of responses in our case studies indicate that linguistic expressions, specific situations of discourse, and particular social and cultural assumptions associated with the expressions are all to be considered as potential factors that, along the constraints imposed by the salience of syntactic position, associated attention state, thematic role, or distance from the antecedent, foregrounded by various extant approaches as discussed in section 2, contribute to the selection of the antecedent. What began as an inquiry into the grammar/pragmatics interface can now be subjected to a search for more precisely formulated factors to do with default scenarios, social and cultural conventions, and, in the absence of these, online pragmatic inference. In short, any further constraints, or, as we decided to call them, factors, will help systematise what on a cursory glance looks simply like contextual contribution.

It is worth elaborating on the fact that one of the case studies (CS3) yielded a preference for S anaphora interpretations across all three conditions. However, as we pointed out above, this is an isolated case among our five studies. We hypothesise that these results can be attributed to the clearly asymmetric and more culturally rigid relationship evoked by the specific lexical items in the example. The phrase 'the president', in conjunction with 'nominate' and 'expect', have arguably led the majority of respondents to prefer subject-bound interpretations and to ignore or reject the possibility of the object-driven reference resolution even in object-biased contexts. Simply put, if A has the asymmetric power to 'nominate' B, in particular if A is 'the president', A can be more frequently taken to 'expect a positive result'. So, in the search for the features for the model of preference orders, we can probably ascribe this result to default scenarios or even the collocations of lexical items (see e.g. Asher 2011) rather than to conventions stemming from grammar and/or information structure if we are aiming at a model that would account for the phenomenon at large rather than at construction-specific preferences. In section 4, we propose that binding preferences can be allocated to different processes identified in DS that are active in discourse interpretations and that they can be mapped onto specific sources of information. In DS-theoretic terms, the feature that is relevant here will pertain to the source dubbed 'stereotypes and presumptions about society and culture', which can converge with the effect of specific linguistic expressions themselves (in DS: the source 'word meaning and sentence structure').

Next, in CS1, and CS2, the out-of-context assessment reveals different preferences (clear object-bias in the first and pretty even distribution in the latter), but both case studies yield easily to contextual biasing. CS2 was particularly conducive to contextual manipulation: 'they' can equally plausibly refer to disappointed students and disappointed teachers, so lexical or conventional bias is not a factor here. Again, an appeal to stereotypes and presumptions about society and culture (or, in terms of processes, automatic interpretations to do with 'socio-cultural and world-

knowledge defaults', see section 4.2) indicates that, arguably, 'congratulate' and 'start a new challenge' in CS1, and 'fail' and 'feel disappointed' in CS2, might have been affected by default scenarios in Q1, albeit to a different degree. If A 'congratulates' or 'fails' B, B is intuitively more likely, in the absence of context, to 'start a new challenge' or to 'feel disappointed' in her failure. Nevertheless, the comparison of responses indicates that these concepts appear less asymmetric and culturally rigid than 'the president', 'nominate', and 'expect' in CS3. It is particularly notable that in CS1, the strong preference for binding to object in the context-free scenario does not yield to further strengthening through contextual biasing to object but remains approximately the same (decreased by one response, or 3.4 per cent). Since this is the object, not the subject position, and the pattern is not repeated in our other case studies, this strongly suggests an explanation through a supplied default scenario.

Finally, CS4 and CS5 are interesting for different reasons: consultants favoured O anaphora in all three conditions and the contextual biasing towards object did not produce the expected raise. Contextual biasing to subject was much more effective than biasing to object. We conjecture that CS4 triggered a presumption that the person was brought to an event for a salient reason. Likewise, CS5 triggered a presumption that the person was invited to a party with the expectation of enjoying it. Because the contexts provided for CS4 and CS5 were worded in exactly the opposite direction, we conclude that the particular lexical items involved, namely 'bring' in CS4 and 'invite' in CS5, triggered the scenarios responsible for consultants' O preferences. So, we have a case here of the lexical contribution to the typology of sources of preference in anaphora resolution. To generalise somewhat, it is plausible to stipulate that in all but the third case study, linguistic expressions might have some effect on the bias towards binding to object. It might also be the case, based on our observations in this section, that 'feel' (CS2) is less object-biased than 'start' (CS1), 'bring' (CS4) or 'invite' (CS5). But, to repeat, our aim was not to determine the exact extent to which lexical items affect anaphora resolution; our aim was merely to look into the possibility of a typology of the sources of such preferences in view of incorporating the effects of a suitable variety of factors in the semantic *qua* conceptual representation.

Although care was taken not to introduce bias, CS4 and CS5 appear to behave analogous to Winograd's seminal examples in (3a,b) discussed in section 2.1, albeit with case-specific strength – and, more importantly, strength that is gauged differently by different consultants. What is of particular interest here is that in the case of CS4, the context-free scenario resulted in a substantially higher proportion of O responses than the contextually object-biased scenario. This seems to show that a tacit default scenario supplied by the consultants in Q1 had a stronger effect than the biasing context in Q3 that would require more elaborate inferencing – and that, perhaps, even had a weakening effect on the conventional scenario. The context supplied in Q3 and the relevant pair of sentences are repeated in (18).

(18) [Kate does not like to support charities. Unlike Lilian, Kate is not generous at all.]

'Kate brought Lilian to the meeting. She made a donation.'

Bearing in mind that in Q2 and Q3 every effort was taken to provide maximally biasing (while at the same naturally sounding) contexts for all items alike, this, again, points towards a need for a model that would incorporate information from various linguistic and non-linguistic sources – the lexicon, as well as socio-cultural conventions and inferences.

But it has to be pointed out that the methodology employed here has some (unavoidable it seems) limitations. The rationale behind CS1-CS5 was to produce context-free, potentially referentially ambiguous examples of cross-sentential pronominal anaphora. However, the availability of two plausible binding sites is not the same thing as absolutely neutral, non-biased constructions. The latter is not attainable in principle in that, as we have indicated, utterance interpretation comes with the construction of salient, default context – and, indeed, we had to encourage consultants to think of scenarios that spring to mind on reading the pairs of sentences in order to elicit the required information concerning binding. In naturally occurring discourse, where prior context is not provided, the interpreter is always likely to supply an imaginary one, drawing on experiences, socio-cultural background assumptions, knowledge of the lexicon including standard collocations, and so forth. So, even when inference from context is not possible, conventions, assumptions, and various kinds of knowledge (of culture, society, human nature, natural laws, and so on) can be drawn on.[18] This is a fact of conversation but this fact also makes us wary of a potential circularity: we employed scenarios that were by necessity not completely unbiased and we drew conclusions concerning contextual bias by using contextualised, biased counterparts (Q2 and Q3). This potential circularity was avoided here in two different ways. First, the degree of neutrality in Q1 constituted a research question in itself and obtained an answer in each of the five case studies through the quantitative results. What followed was the contrast between the behaviour of the expressions in Q1 assessed not as a study in a neutral scenario but precisely, and correctly, in a scenario where the only bias could come from the preferences dictated by a salient imagined scenario to which socio-cultural conventions, lexical content, personal experiences, and other factors contributed. Next, and as a result, we compared the manipulated biases in Q2 and Q3 with what we *actually found* through Q1, not with the idealisation of a neutral, unbiased scenario. In other words, the results of Q1 were factored into the two-way contextual shifts.

Such a construal allows us to proceed from the analysis of the found preferences to the stipulated factors. It also corroborates our earlier scepticism about large-scale quantitative studies of the phenomenon that aim to test more than specific constructions in that quantifying over large data would smooth out what is attributable to different causes. We believe that if more extensive experimental sup-

[18] Our instructions in Q1 exploited this fact.

port is to be sought, it ought to be a post-theoretic step and consist of a scrutiny of large corpora in order to address specific questions that the theory throws up, such as associating the type of resolution with the type of the source and processes the theory identifies or, as a more promising route, neuroimaging that shows the areas and patterns of activation which can then be mapped onto such interacting processes. We will have more to say about sources of information and processes that contribute to a conceptual representation in section 4.

In short, the results of our five case studies show that, if one aims at a comprehensive model of cross-sentential anaphora resolution, the diversity of factors that affect cross-sentential pronominal anaphora resolution is indeed to be embraced rather than fit into tight rules. The results strongly point towards a model of cross-sentential pronominal anaphora resolution that makes use of various factors to which we referred in our discussion – and more. Amalgamating all current and previous findings allows us to proceed to proposing a preliminary typology of such factors and subsequently to building conceptual-semantic representations. Only when we adopt a very different outlook that allows us to see coherence, attention, and structure not as competitors but as contributors to the relevant processes, acting on a par with other key factors, do we stand a chance of constructing a model and, possibly, an associated normative theory of anaphora resolution as part of a theory of discourse interpretation.

If one is willing to continue on this methodologically 'positively eclectic' route and hunt and gather the lessons from the subsequent DR-theoretic developments, Kamp and Reyle's 'extraneous clues' begin to translate into an array of factors. All in all, our inquiry revealed the pertinence of at least the following factors:

(i) socio-cultural stereotypes and conventions,
(ii) typical scenarios encountered in conversation,
(iii) background general knowledge,
(iv) distance of the antecedent along the memory line,
(v) lexical content and inferences form lexical meaning,
(vi) sentence structure,
(vii) inference from context,

as well as

(viii) relative strength of a factor *vis-à-vis* other active factors (for example the strength of the contextual effect on the interpretation *vis-à-vis* the stereotypical scenarios).

To this, we can add a rather uncontroversial potential factor.

(ix) logical inference, both non-defeasible (deduction) and defeasible (induction and abduction).

Finally, we add what is rarely remembered about:

(x) individual experiences of the interpreter that have the power to affect the interpretation.

Just as pointing and other gestures are on a par with linguistic anaphora and has to be accounted for in a semantic *qua* conceptual representation, so do other more nuanced factors such as personal preferences and experience that can give rise to variable activation of referents.[19]

This tentative list has no claim to exhaustivity; neither would it benefit from a search for it. What we want to flag is that such factors have their place in a conceptual-semantic representation of cross-sentential anaphora – at least of the pronominal kind, if not at large – and that, *a fortiori*, the question as to whether one ought to place conventions in the grammar, in the middle level of 'glue logic', or entirely in pragmatics, does not have much bearing on the representation of discourse: the relevant factors are all there in their own right and their contribution ought to be represented if our aim is a semantic representation with a claim to the status of mental representation. So, the next step is to allocate the relevant interpretive processes to the above factors and proceed to the representations. We attempt it in section 4.

4 Conceptual-Semantic Structure and Pragmatic Composition: Towards a new model of cross-sentential pronominal anaphora resolution

4.1 Sources, processes and factors: representing cross-sentential anaphora in DS

The question that we have to address now is how to transform this eclectic collection of relevant factors and observations gathered from DR-theoretic, coherence-theoretic and attention-based approaches, as well as from our five case studies, into a model of a mechanism of cross-sentential pronominal anaphora resolution that would enjoy a better predictive power. As was briefly anticipated in section 3, we propose allocating the factors to sources of information that interactants utilise in discourse processing. DS (e.g. Jaszczolt 2005, 2010) recognises five such sources: (i) word meaning and sentence structure (WS), that is the lexicon and grammar that produce standard logical forms; (ii) world knowledge (WK) that can lead to inferences from, say, general physical laws to the interpretation of particular situations or to the employment of default reasoning from such laws; (iii) situation of discourse that facilitates inferring (SD); (iv) properties of human inferential system (IS), responsible crucially for the presumption of maximal informativeness of utterances (and the associated maximal intentionality of mental states), such as for example referential rather than attributive, presupposing rather than non-presupposing, or *de re* rather than *de dicto* interpretation; and (v) stereotypes and

[19] See here Giora's (2003) graded salience hypothesis. Her experiments suggest that meaning activation is dictated by factors such as frequency, familiarity, conventionality and prototypicality and as such (i) varies with individual experience and (ii) is not entirely dictated by contextual preferences: irrelevant interpretations that are nevertheless salient, 'foremost on our mind' (p.10) can attain a certain level of activation.

presumptions about society and culture (SC) that can be resorted to in default reasoning or in inferential processes.[20]. These sources provide information that contributes to the conceptual representation on an equal footing. As a result, we avoid the largely metatheoretic questions concerning the status of regularities as linguistic or non-linguistic rules. Naturally, the sources of information could have been distinguished with a different degree of granularity but, since the DS-theoretic way of distinguishing them has proven successful in its implementations for the analysis of different discourse phenomena in a range of languages, let us take them as a matrix for our allocation of factors, bearing this degree of arbitrariness in mind.

First, the projection line for anaphora resolution, in unambiguous cases (that is in cases where there is only one suitable antecedent) draws on the lexicon and grammar, and therefore will be allocated to WS. Van der Sandt's constraint from partial matches[21] also falls under this source, as fully matching nominal phrases are preferred and this information can be found in the sentence itself. Next, as discussed above, conventions that stem out of coherence will enter our model in the form of default scenarios.[22] So, these will be allocated to WK and SC. Third, IS contributes the maximisation of informative content – a feature of discourse amply discussed in post-Gricean pragmatics and variously allocated to post-Gricean principles and heuristics. DS places this feature in the IS source and ultimately in the strength of intentionality of the associated mental states. Finally, SD captures other attested anaphoric connections that do not fall under the above generalisations but stem out of the current situational context. This completes the palette of the players to be used in our model – and in accounting for the diversified findings in our case studies. The results of these studies are fully compatible with the DS-theoretic stance that information from the WS source interacts with information from SD, IS SC, and WK in producing semantic representations. Applied to the phenomenon of cross-sentential pronominal anaphora, this suggests that a model of reference assignment in such constructions can be constructed utilising the sources and the processes of interpretation associated with them – as we demonstrate in what follows.

[20] For a more detailed introduction to sources of information and the associated processes used in the derivation of an interpretation in DS see e.g. Jaszczolt (2010). For some applications see e.g. Jaszczolt (2005, 2016a), or Elder (2019).

[21] See footnote 9.

[22] The term 'conventions' is used in the literature in a variety of ways. Conventions can be understood as standards that are difficult to override or as mere precedent scenarios that result in slightly greater salience than that of alternative interpretations. Geurts (2018: 118) argues, for example, that conventions are 'precedent types' enabled by common ground and they can be 'lawlike' (like conventions of grammar) or 'merely enabling,' like preferred interpretations. Following this distinction, conventions that lexical items and grammatical structures give rise to are 'merely enabling,' so we could in principle also represent them as SCWD-triggered. But because they are triggered by linguistic structures rather than precedents in the sense of past scenarios, we have opted for including them under WS. As Geurts (2018: 120-121) argues, although words and syntactic categories are not themselves conventional, they individuate conventional speech actions and as such are 'conventionally used.'

4.2 Merger representations

We now provide contextualist semantic *qua* conceptual representations of cross-sentential pronominal anaphora that capture the fact that a combination of linguistic and non-linguistic sources of information leads language users to either subject- or object-bound resolutions. The contextualist framework of DS we are employing is an offshoot of DRT that utilises an amended and extended language of DRSs but aims at representing a somewhat different, more 'pragmaticky' object from that represented in DRT in that it represents primary acts of communication: primary meanings, intended and recovered as such by interlocutors in situations where conversation proceeds according to mutually accepted (Gricean) principles and heuristics. Like DRSs, its representations (called *merger representations*) have the status of mental representations, but unlike DRSs they represent speech act content.[23] As such, they can also represent content that is indirectly conveyed. This is possible thanks to the abandonment of what DS theory calls the syntactic constraint: unlike in other DR-theoretic accounts (and, for that matter, also post-Gricean truth-conditional accounts of meaning), the logical form, the output of syntactic processing of the sentence, is not regarded as a starting point for the construction of the semantic representation, only to be modified, developed, enriched, modulated, and so forth (depending on the metalanguage of a particular theory). Instead, all linguistic and non-linguistic sources of information about meaning discussed above in section 4.2 contribute to the representation on an equal footing. So, it follows that in the case of indirect communication, the result of inference or an adoption of a default interpretation arrived at through conventions or stereotypes can 'override', so to speak, the output of syntactic processing: it is the primary intended meaning that is modelled, not the contextually embellished sentence structure. For example, unlike in DRSs, tense-time mismatches are easily explicable in terms of such defaults or pragmatic inference.

Merger representations are truth-conditional representations. Compositionality is understood in DS as pragmatic, interactive compositionality that obtains on the level of mental representations. To repeat, within the DS framework, Jaszczolt (2005, 2010, 2016a) purports that the outputs of five linguistic and non-linguistic sources of information merge to yield speakers' main messages, or primary meanings, that hearers successfully recover. These five *sources* map, in different contexts, onto four different *processing mechanisms*: WS maps onto processing that is specific to WS;[24] WK and SC map onto either social, cultural, and world knowledge defaults (SCWD) or conscious pragmatic inference (CPI); SD maps onto CPI; and IS maps onto cognitive defaults (CD). Merger representations provide information

[23] DS can be called a radically contextualist approach in the sense of representing primary intended meanings irrespective of their method of conveyance. However, its stance on Wittgenstein-Travis occasionalism, that is concerning the question of the context-dependence of word meaning in the WS source/process, is still the subject of research in progress. For now, suffice it to say that the lexicon will be taken to contain conceptual cores and a pragmatic overlay. For different senses in which a contextualist stance can be 'radical' see Jaszczolt (in press).

[24] Jaszczolt (2016a) leaves the question of modularity open and argues that DS as it stands is compatible both with the modularity assumption and with general processing mechanisms.

about these processes and point out the material on which each process operates, so our earlier discussion of the factors encountered in anaphora resolution will now lead to representations in terms of types of processes and the pertinent material. The correlation with the *factors* in our attempted typology (in section 3.4) is as follows: (i) and (ii) will correspond to the process SCWD; (iii) to CPI or SCWD; (iv) and (ix) to CD; (v) and (vi) to WS; (vii) and (x) to CPI; (viii) is a relative concept. Since DS is by now a well-established framework, in what follows we solely explain the notions that are pertinent for our specific analyses.

Let us first consider anaphora in CS5, repeated in (19), that obtained 73.3 per cent of O interpretations in the context-free environment.

(19) Tom invited Andrew. He enjoyed the party a lot.

If (19) is presented out of context, the consultants can be taken to make use of the factors (i) and (ii), as well as the 'ubiquitous' factors (viii - x). Factors (v) and (vi) can be excluded as they would have been likely to have led to uniform results and to a greater resistance to contextual bias in Q2 and Q3. However, as was remarked in an earlier disclaimer, since the lexical content itself leads to lexicon- (rather than encyclopaedia-) based inferences, the boundary between (i - iii) on the one hand and (v) on the other is quite theory-dependent. In other words, we could allocate the choice of the antecedent to lexical inference that inviting a guest comes with a wish to make them enjoy the event. But, using common sense judgement, it appears more appropriate to allocate this information to stereotypes, customs and conventions pertaining to culture and society rather than to lexical knowledge.

To repeat, in DS-theoretic terms, factors (i) and (ii) pertain to the process SCWD. Factor (iv) that is potentially relevant pertains to CD in that it draws on the capacity of the working memory of the interpreting agent. On this scenario, however, it is not likely that this feature would normally be invoked. Figure 2 shows the merger representation corresponding to the SCWD-driven interpretation of the anaphora in (19), where 'Tom' and 'Andrew' are processed via CD (in virtue of being directly referential expressions) and the primary meaning is communicated directly, so it overtly follows the logical form of the sentences (WS).[25]

Next, the preferred S interpretation of example (1), or our CS3, repeated below, can also be attributed to the SCWD-induced anaphoric binding in the context-free scenario in which this reading featured in 76.7 per cent of responses.

(1) The president nominated Jones. He expected a quick confirmation.

In (1), social and cultural assumptions could be taken to be, say, that presidents, being in a position of power, nominate less powerful individuals, and that it is presidents who are usually in the position to expect a quick confirmation. At the same time, as we argued above, the lexical items 'president' and 'nominate' strongly

[25] For the sake of brevity, our merger representations introduce only the discourse referents that are relevant for our discussion of anaphora; 'the party' in Figure 2, for example, introduces its own discourse referent (say, z). For the same reason, the representation of temporal reference is omitted here. For representing time in DS see e.g. Jaszczolt (2009).

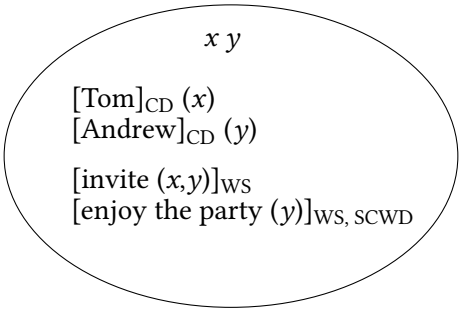

Figure 2 Merger representation for an object-bound, context-free interpretation of (19).

suggest that 'expecting confirmation' is to be attributed to the president. The entrenchment of this interpretation in the adverse contextual bias also suggests the importance of the feature (v) (lexical content and inferences from lexical meaning). The representation is provided in Figure 3.

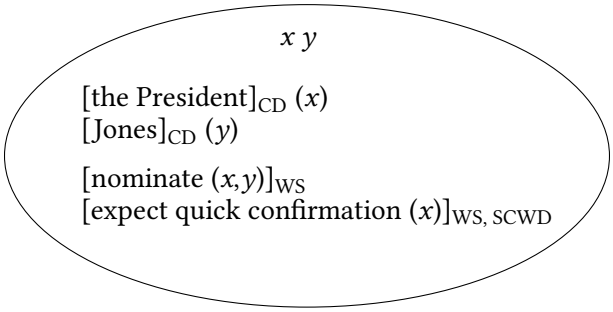

Figure 3 Merger representation for a subject-bound, context-free interpretation of (1).

The fact that the anaphora resolution is here attributed partly to WS does not, however, constitute support for a grammar-based account *à la* Lepore & Stone (2015). The reading is attributed here to the concepts associated with the particular lexical items in conjunction with social conventions: pinning it on WS alone would make it difficult to account for cases of alternative binding.

Next, cross-sentential pronominal anaphora in contextualised environments can be represented as obtained via CPI or SCWD, or even WS alone when (a) the binding site does not change as compared with the context-free scenario, (b) binding was triggered by the lexicon or grammar in the first place, and (c) the context did not show significant effects. We will now consider the situation in CS3 where the context did alter the binding preference. If it is clear from context that Jones was in desperate need of employment (and that the president was not in a hurry to employ him), CPI-driven inferential O interpretation can be favoured. This reading is represented in Figure 4.

Other factors that affect anaphora resolution can be represented analogously, in virtue of the correlations established here. Needless to say, neither the typology of factors nor the inventory of processes offered by the theory of DS constitute the

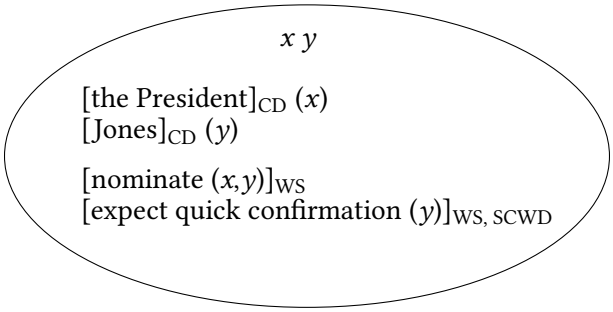

Figure 4 Merger representation for an object-bound interpretation of (1) in an object-biasing context.

final word on this important issue. But they are a step towards the recognition of its complexity.

5 Conclusions, Disclaimers, and Future Predictions

In this paper we presented a sketch of a multidimensional, pragmatics-rich account of cross-sentential pronominal anaphora. We argued that there is nothing in the grammar of English that comprehensively leads to either subject- or object-biased reference assignments. Our theoretical inquiry, paired with some insight from our five case studies, point towards the importance of a set of factors that jointly delimit the intended interpretation. We tentatively listed them as (i - x), pointing out that the search for a precise typology is neither necessary for our objective at hand, nor is it possible without a large empirical study, either based on large corpora or on neuroimaging, that would correlate instances of reference assignments with sets of interacting factors and, ideally, also with a value standing for their relative strength. But this is a separate task for the future that is best left to computational linguistics or to neuropragmatics. We suggested that such empirical work could only be undertaken post-theoretically, when an informed assumption can be made as to what the relevant correlates are and, *a fortiori*, what correlations to look for. Next, with the help of the representations in the contextualist theory of DS, we illustrated how different processes used in the derivation yield the primary meanings that speakers intend to communicate and that hearers successfully recover. This was possible because semantic representations are considered in DS to be conceptual representations; they capture the primary intended meaning independently of the sources or information and processes used for its conveyance. We have argued that when compositionality is to be sought at the level of conceptual structure, the theory of anaphora resolution can enjoy greater normativity and, as such, greater predictive power. While linguistic structures and their logical form constitute a valid object of philosophical and linguistic analysis, it appears that once we address the question of speaker's choices in meaning assignment, meaning representations have to concern themselves with many more sources of information and many more processes than the composition of sentence meaning alone. This desideratum openly voiced in the DR-theoretic accounts discussed in section 2 was taken very seriously in our

current study and led to the adoption, following DS, of a pragmatics-rich object of truth-conditional analysis in the form of acts of communication and their merger representations.

In the context of the ongoing debates between different approaches to pronominal anaphora resolution, including debates between some seminal ones discussed here, it appears that, instead of opting for one of these orientations or any local divisions of the field, it would be more prudent to posit that two kinds of mechanisms interact in the process: (i) general cognitive mechanisms account for pragmatic inference and for resolution based on socio-cultural preferences and the resulting salience that inform coherence theory, and (ii) language-specific mechanisms pertaining to information structure, focused on, most notably, by Centering Theory, account for preferences associated with positions of focus and as such with activation, attention state, and in general, the cognitive status. As we observed in the discussion of contextual bias in our five case studies, it would seem plausible to interpret the differences between the subject- and object-bias by attributing the first to the salient position in information structure while the latter to the preferences derived from the lexical meaning and discourse coherence. However, such straightforward eclecticism did not appear warranted. We also observed the input from lexical and pragmatic information and from bringing to salience. The latter can foreground different grammatical positions – and this proves to be, quite often, the object position, the position that in the typical structure of an English sentence is less distant from the antecedent. So, if there are structure-based preferences, they appear to be more complex than the ones predicted by attention theorists. Instead, the structural proximity of the object and the often focal position of the subject can compete for preference. But it also has to be noted that in our typology of factors set out in section 3.4 we captured the distance of the object in terms of a distance on the memory line (factor (iv)), in that the latter will often coincide with the first. What appears to be a structural factor can also be given a psycholinguistic interpretation. The need for positing such a typology, as well as certain degree of arbitrariness inherent in any such attempt, allowed us to conclude that the inquiry into the grammar/pragmatic interface in the context of cross-sentential pronominal anaphora resolution largely loses its appeal. As we put it, the diversity has to be embraced rather than squeezed into ill-fitting regularities.

Finally, in addition to the diversity of factors, sources and processes that have to play a part in a theory of anaphora resolution, there is another kind of diversity that requires an acknowledgement. Our discussion in this paper spans theories that emerge out of computational approaches to language, as well as Gricean theories that, on the whole, subscribe to a greater dose of psychologism, evident in the intention- (and intention recognition-) driven object of study, or emphasis on cooperativity of conversation. These different aims are not irreconcilable but they can create difficulties. In their well-quoted and somewhat contentious book that inspired this article, Lepore & Stone (2015: 83) defend the view that '[p]ragmatics can be, at most, a theory of disambiguation; pragmatic reasoning never contributes content to utterances'. We hope that our inquiry has demonstrated that pragmatics does much more, and at the same time much less, than that. First, a variety

of pragmatic factors identified here contribute to discourse anaphora resolution in specific ways. On the other hand, pragmatic factors normally are not in the business of disambiguating in that, on standard scenarios when background assumptions are correctly gauged, there is no ambiguity to resolve. There is the speaker's meaning to recover, using the available processes and sources at hand, where the grammar is merely one of several players. Ambiguity in discourse is bad practice, unless it serves the purpose of *un jeu de mots*. Likewise, postulating ambiguity in the analysis of discourse would be bad practice if it had no external justification. In computational linguistics, this justification is provided by application to machine learning. But, needless to say, Gricean pragmatics does not abide such a justification but rather follows the methodological directive specified in Modified Occam's Razor not to proliferate senses where different uses can be accounted for by pragmatic principles. This, of course, does not mean that the ball is in the Gricean court. It means that there are different games to be played, one of which focuses on intentions and intention recovery in human communication, while the other on rules that lead to the inventory of meanings that humans and machines likewise can convey with the help of language systems, conventions, and contexts. In the latter, computational linguists' game, ambiguity can remain rife and a search for a powerful concept of 'grammar that does it all' is a justifiable search. We hope that different aspects of our proposal will inform both games.

References

Arundale, R. B. 1999. An Alternative Model and Ideology of Communication for an Alternative to Politeness Theory. *Pragmatics* 9(1). 119–153.

Arundale, R. B. 2010. Constituting Face in Conversation: Face, facework, and interactional achievement. *Journal of Pragmatics* 42(8). 2078–2105.

Asher, N. 2011. *Lexical Meaning in Context: A web of words*. Cambridge: Cambridge University Press.

Asher, N. & A. Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.

Blutner, R. & H. Zeevat (eds.). 2004. *Optimality Theory and Pragmatics*. Basingstoke: Palgrave Macmillan.

Dekker, P. J. E. 2012. *Dynamic Semantics*. Dordrecht: Springer.

Elder, C.-H. 2019. *Context, Cognition and Conditionals*. Basingstoke: Palgrave Macmillan.

Ferretti, T. R., H. Rohde, A. Kehler & M. Crutchley. 2009. Verb Aspect, Event Structure, and Coreferential Processing. *Journal of Memory and Language* 61(2). 191–205.

Geurts, B. 2018. Convention and Common Ground. *Mind & Language* 33(2). 115–129.

Giora, R. 2003. *On Our Mind: Salience, context, and figurative language*. Oxford: Oxford University Press.

Grice, P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Grosz, B. J., A. K. Joshi & S. Weinstein. 1995. Centering: A framework for modelling the local coherence of discourse. *Computational Linguistics* 21(2). 203–225.

Grosz, B. J. & C. L. Sidner. 1998. Lost Intuitions and Forgotten Intentions. In M. A. Walker, A. K. Joshi & E. F. Prince (eds.), *Centering Theory in Discourse*, 39–51. Oxford: Clarendon Press.

Gundel, J. K. 1998. Centering Theory and the Gvenness Hierarchy: Towards a synthesis. In M. A. Walker, A. K. Joshi & E. F. Prince (eds.), *Centering Theory in Discourse*, 183–198. Oxford: Clarendon Press.

Gundel, J. K. 1999. On Different Kinds of Focus. In P. Bosch & R. van der Sandt (eds.), *Focus: Linguistic, cognitive, and computational perspectives*, 293–305. Cambridge: Cambridge University Press.

Gundel, J. K., N. Hedberg & R. Zacharski. 1993. Cognitive Status and the Form of Referring Expressions in Discourse. *Language* 69(2). 274–307.

Hobbs, J. R. 1979. Coherence and Coreference. *Cognitive Science* 3(1). 67–90.

Hobbs, J. R. 1990. *Literature and Cognition.* Stanford, CA: CSLI.

Jaszczolt, K. M. 2005. *Default Semantics: Foundations of a compositional theory of acts of communication.* Oxford: Oxford University Press.

Jaszczolt, K. M. 2008. Psychological Explanations in Gricean Pragmatics and Frege's Legacy. In I. Kecskes & J. Mey (eds.), *Intentions, Common Ground, and Egocentric Speaker-Hearer*, 9–44. Berlin: Mouton de Gruyter.

Jaszczolt, K. M. 2009. *Representing Time: An essay on temporality as modality.* Oxford: Oxford University Press.

Jaszczolt, K. M. 2010. Default Semantics. In H. Heine & H. Narrog (eds.), *The Oxford Handbook of Linguistic Analysis*, 193–221. Oxford: Oxford University Press.

Jaszczolt, K. M. 2016a. *Meaning in Linguistic Interaction: Semantics, metasemantics, philosophy of language.* Oxford: Oxford University Press.

Jaszczolt, K. M. 2016b. On Unimaginative Imagination and Conventional Conventions: A response to Lepore and Stone. *Polish Journal of Philosophy* 10(1). 89–98.

Jaszczolt, K. M. 2019. Rethinking being Gricean: New challenges for metapragmatics. *Journal of Pragmatics* 145. 15–24.

Jaszczolt, K. M. in press. Post-Gricean Pragmatics for Intercultural Communication. In I. Kecskes (ed.), *The Cambridge Handbook of Intercultural Pragmatics*, Cambridge: Cambridge University Press.

Jaszczolt, K. M. & R. Sileo. 2021. Pragmatics and Grammar as Sources of Temporal Ordering in Discourse: The case of *and.* In F. Macagno & A. Capone (eds.), *Inquiries in Philosophical Pragmatics: Issues in Linguistics*, 53–81. Cham: Springer.

Kamp, H. & U. Reyle. 1993. *From Discourse to Logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory.* Dordrecht: Kluwer.

Kehler, A. 2002. *Coherence, Reference, and the Theory of Grammar.* Stanford, CA: CSLI.

Kehler, A., L. Kertz, H. Rohde & J. L. Elman. 2008. Coherence and Coreference Revisited. *Journal of Semantics* 25(1). 1–44.

Kehler, A. & H. Rohde. 2013. A Probabilistic Reconciliation of Coherence-Driven and Centering-Driven Theories of Pronoun Interpretation. *Theoretical Linguis-*

tics 39(1-2). 1–37.
Lepore, E. & M. Stone. 2015. *Imagination and Convention: Distinguishing grammar and inference in language.* Oxford: Oxford University Press.
Levinson, S. C. 2000. *Presumptive Meanings: The theory of generalized conversational implicature.* Cambridge, MA: MIT Press.
Lewis, D. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic* 8(1). 339–359.
Parikh, P. 2016. Comments on *Meaning in Linguistic Interaction.* Presented at Book Symposium on Kasia M. Jaszczolt *Meaning in Linguistic Interaction*, American Philosophical Association Pacific Division Meeting, San Francisco. Acessed 3 June 2020. https://www.academia.edu/24318355/P._Parikh_Comments_on_K._Jaszczolts_Meaning_in_Linguistic_Interaction_APA_Pacific_San_Francisco_2016_.
van der Sandt, R. A. 1992. Presupposition Projection as Anaphora Resolution. *Journal of Semantics* 9(4). 333–377.
van der Sandt, R. A. 2012. Presupposition and Accommodation Discourse. In K. Allan & K. M. Jaszczolt (eds.), *The Cambridge Handbook of Pragmatics*, 329–350. Cambridge: Cambridge University Press.
Stojnić, U., M. Stone & E. Lepore. 2017. Discourse and Logical Form: Pronouns, attention and coherence. *Linguistics and Philosophy* 40(5). 519–547.
Stojnić, U., M. Stone & E. Lepore. 2020. Pointing Things Out: In defense of attention and coherence. *Linguistics and Philosophy* 43(2). 139–148.
Strawson, P. F. 1964. Intention and Convention in Speech Acts. *Philosophical Review* 73(4). 439–460.
Travis, C. 2006. Psychologism. In E. Lepore & B. C. Smith (eds.), *The Oxford Handbook of Philosophy of Language*, 103–126. Oxford: Clarendon Press.
Walker, M. A., A. K. Josh & E. F. Prince (eds.). 1998. *Centering Theory in Discourse.* Oxford: Clarendon Press.
Winograd, T. 1972. Understanding Natural Language. *Cognitive Psychology* 3(1). 1–191.

Roberto B. Sileo
The University of Cambridge
rbs30@cam.ac.uk

Kasia M. Jaszczolt
The University of Cambridge
kmj21@cam.ac.uk

Appendix

Q1 (context-free)	**Q2 (subject-biased)**	**Q3 (object-biased)**
CS1	[The grandmother helped the granddaughter to finish her bachelor's degree. The granddaughter wanted to be a student forever.]	[The granddaughter helped the grandmother to finish her bachelor's degree. The grandmother wanted to be a student forever.]
'The granddaughter congratulated the grandmother. She now started a new challenge.'	**'The granddaughter congratulated the grandmother. She now started a new challenge'**	**'The granddaughter congratulated the grandmother. She now started a new challenge.**
CS2	[The students did not care about their results. The teachers had worked very hard and even cancelled their holidays to prepare extra lessons.]	[The teachers did not care about the students. The students had big dreams to go to university.]
'The teachers failed the students. They felt disappointed.'	**'The teachers failed the students. They felt disappointed.'**	**'The teachers failed the students. They felt disappointed.'**
CS2	[Jones was not in a hurry to start working. The president needed a new employee as a matter of urgency.]	[The president really wanted to delay the new employee's starting date. Jones needed to earn money as a matter of urgency.]
'The president nominated Jones. He expected a quick confirmation.'	**'The president nominated Jones. He expected a quick confirmation.'**	**'The president nominated Jones. He expected a quick confirmation.'**
CS4	[Lilian does not like to support charities. Unlike Kate, Lilian is not generous at all.]	[Kate does not like to support charities. Unlike Lilian, Kate is not generous at all.]
'Kate brought Lilian to the meeting. She made a donation.'	**'Kate brought Lilian to the meeting. She made a donation.'**	**'Kate brought Lilian to the meeting. She made a donation.'**
CS5	[Andrew did not know how to behave in company and so, he was not Tom's favourite person. On this occasion, Andrew did not do anything crazy at the party, though.]	[Tom did not know how to behave in company and so, he was not Andrew's favourite person. On this occasion, Tom did not do anything crazy at the party, though.]
'Tom invited Andrew. He enjoyed the party a lot.'	**'Tom invited Andrew. He enjoyed the party a lot.'**	**'Tom invited Andrew. He enjoyed the party a lot.'**

Table 1 Pairs of sentences and contexts in Q1-Q3.

Q1												
Responses	**CS1**	**%**	**CS2**	**%**	**CS3**	**%**	**CS4**	**%**	**CS5**	**%**	**Total**	**%**
S	4	13.3%	12	40.0%	23	76.7%	3	10.0%	1	3.3%	43	28.7%
0	26	86.7%	17	56.7%	5	16.7%	24	80.0%	22	73.3%	94	62.7%
Total S & O*	30	100.0%	29	96.7%	28	93.3%	27	90.0%	23	76.7%	137	91.3%
N/C	0	0.0%	1	3.3%	2	6.7%	3	10.0%	7	23.3%	13	8.7%
Total	30	100%	30	100%	30	100%	30	100%	30	100%	150	100%

Q2												
Responses	**CS1**	**%**	**CS2**	**%**	**CS3**	**%**	**CS4**	**%**	**CS5**	**%**	**Total**	**%**
S	22	73.3%	26	86.7%	28	93.3%	10	33.3%	10	33.3%	96	64.0%
O	7	23.3%	4	13.3%	1	3.3%	19	63.3%	11	36.7%	42	28.0%
Total S & O*	29	96.7%	30	100.0%	29	96.7%	29	96.7%	21	70.0%	138	92.0%
N/C	1	3.3%	0	0.0%	1	3.3%	1	3.3%	9	30.0%	12	8.0%
Total	30	100%	30	100%	30	100%	30	100%	30	100%	150	100%

Q3												
Responses	**CS1**	**%**	**CS2**	**%**	**CS3**	**%**	**CS4**	**%**	**CS5**	**%**	**Total**	**%**
S	4	13.3%	4	13.3%	18	60.0%	13	43.4%	4	13.3%	43	28.7%
O	25	83.3%	26	86.7%	11	36.7%	17	56.7%	18	60.0%	97	64.7%
Total S & O*	29	96.7%	30	100.0%	29	96.7%	30	100.0%	22	73.3%	140	93.3%
N/C	1	3.3%	0	0.0%	1	3.3%	0	0.0%	8	26.7%	10	6.7%
Total	30	100%	30	100%	30	100%	30	100%	30	100%	150	100%

*Percentages correspondig to S and O responses do not always add up to Total S & O percentages due to rounding

Table 2 Numbers and percentages of subject, object and non-classifiable responses for CS1-CS5 in context free, contextually subject-biased and contextually object-biased scenarios (Q1-Q3).

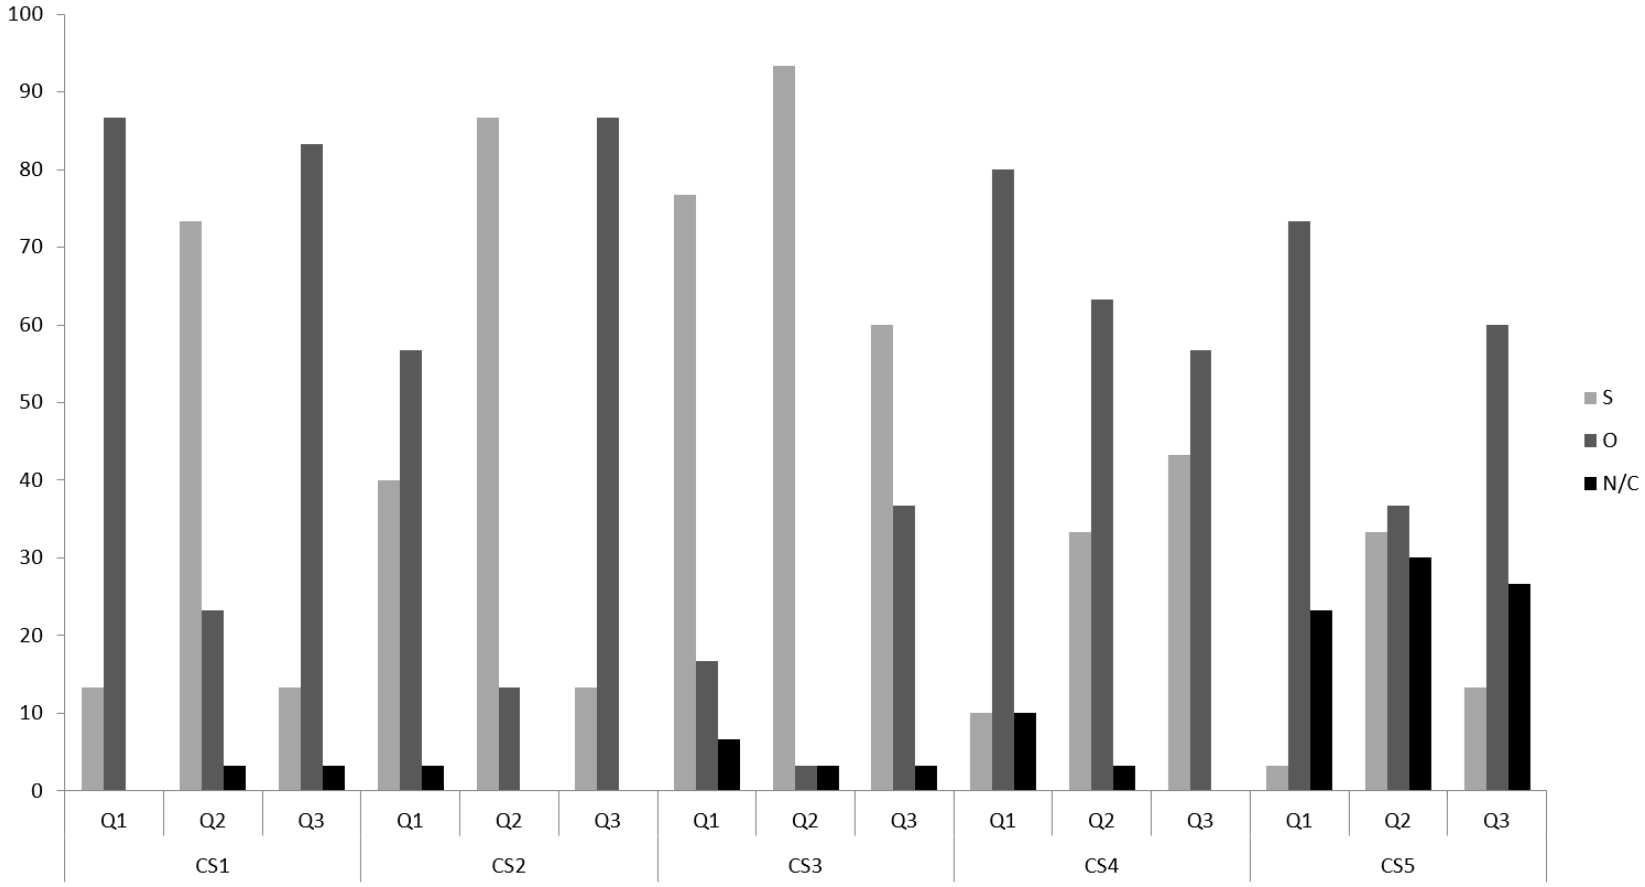

Figure 1 Percentages of subject, object and non-classifiable responses for CS1-CS5 in Q1-Q3.