GLOSSA psycholinguistics

Simons, M., & Rohde, H. (2025). An investigation of definiteness as a trigger of bridging. *Glossa Psycholinguistics*, 4(1): 5, pp. 1–42. DOI: https://doi.org/10.5070/G60112590

eScholarship University of California

An investigation of definiteness as a trigger of bridging

Mandy Simons, Carnegie Mellon University, mandysimons@cmu.edu

Hannah Rohde, University of Edinburgh, hannah.rohde@ed.ac.uk

A *bridged interpretation* of a noun phrase (NP) is one in which the referent is understood to stand in some unstated relation to an entity or event previously mentioned in the discourse. For example, in the sequence *Yasmin approached the house. The door was open.*, the NP *the door* is naturally interpreted as referring to a door of the just-mentioned house. In the theoretical literature, definiteness is often identified as the key driver of bridged interpretations, requiring an alternative analysis for bridged indefinites (*Yasmin approached the house. A door was open.*). We contrast this two-phenomena approach with a one-phenomenon approach, whereby bridging inferences are understood to be the result of general considerations of discourse coherence, particularly facilitated by entity relatedness, but also responsive to effects of definiteness. We present two new methods aimed at measuring the ease and strength of participants' bridging inferences when entity relatedness and definiteness are manipulated. The two-phenomena view predicts that definiteness has a distinctive role to play in inducing bridged interpretations, but contra this view, our results show no independent effect of definiteness. Rather, Experiment 1 (a dialogue-continuation task that probes the presence of bridged interpretations) shows only a main effect of entity relatedness. In Experiment 2 (a self-paced-reading task that probes processing difficulty when a potential bridge is broken), we find an interaction whereby high entity relatedness and the presence of the definite together induce an early commitment to a bridged interpretation. We take these findings to support a unified account in which definite NPs do not require a separate bridging mechanism, but rather are treated like other NPs in being subject to the joint satisfaction of a set of linguistic and more broadly pragmatic constraints.

1. Introduction

When we talk, we often talk about specific people and things, and frequently do so using nouns that apply to multitudes of objects: *(the) house, (a) door, (the) table*. Such general descriptions would plausibly license a listener to ask: "which N?" One thing that obviates the need for this question is the tendency of listeners to infer specific relations between a newly mentioned entity and some entity already under discussion. Thus, hearing utterances of (1a) or (1b), the question, "which door?" doesn't normally arise.

(1) a. Yasmin approached the house. The door was open.
b. Yasmin approached the house. A door was open.

In both cases, the subject noun phrase (NP) of the second sentence, whether definite (*the door*) or indefinite (*a door*), is naturally understood to refer to a door of the just-mentioned house. Indeed, given just this two-sentence sequence, it seems almost impossible to assign this NP any other interpretation.

In other cases, however, an NP may be understood to introduce a new entity not standing in any relevant relation to a previously mentioned entity. Compare (1a) with (2):

(2) Yasmin approached the house. The store was open.

Here, we have no inclination to posit any particular relation, or any relation at all, between the newly mentioned store and the previously mentioned house. Instead, we assume that the NP *the store* refers to some store that is supposed to be identifiable by the hearer/reader (as with the NP *the house* in the first sentence).

We will call the interpretation of *the/a door* in (1) a *bridged* interpretation. In a bridged interpretation, as we will understand it here, the interpreter infers a relation between a newly mentioned entity and some entity or event currently in their discourse model, a relation not made explicit in the discourse; what we have in mind is what Hawkins (1978) calls an associative anaphoric interpretation. We will discuss this way of defining bridging and compare it with other characterizations in Section 2.

Bridged interpretations of NPs are common and seem to present no interpretative difficulties; but clearly not all NPs can be bridged. Given the ubiquity of bridging, an important question for both theoretical and cognitive approaches to discourse interpretation is what triggers the construction of a bridged interpretation, and what prevents such an interpretation from arising (as in (2)). We focus on the following question: is bridging of definites and of indefinites triggered in the same way? That is, are examples (1a) and (1b) above instances of a single phenomenon of bridging, or are they, despite their apparent similarity, actually instances of two distinct phenomena? To explain why we focus on this question, we need to set it in the context of prior work on bridging and prior assumptions about which data is relevant for analysis.

The phenomenon of bridging enters the theoretical semantics/pragmatics literature in relation to definite NPs (henceforward, simply *definites*), as in (1a). Clark (1975) introduced the term *bridging* to characterize a particular set of inferences which, in his account, are a consequence of Givenness marking. These inferences include those which support bridging of definites. As indefinite NPs (henceforward, *indefinites*) are not marked as Given, Clark's account does not extend to cases such as (1b) above.[1] (We will discuss Clark's account in more depth in Section 2.) The extensive literature on definiteness has continued to theorize about bridged interpretations of definites, focusing on how bridging can serve to satisfy the constraints on the use of these expressions. Indeed, in that literature, it is widely argued that bridged interpretations of definites are driven by the need to satisfy those constraints. For some theorists, the term *bridging* just refers to the satisfaction of the requirements of a definite by constructing an appropriate relation.[2] Hence, that literature sees bridging as an interpretational effect triggered by the presence of definiteness marking. On this understanding of bridging, example (1b), as per Clark, is simply not an instance of the same phenomenon. This may seem like a terminological debate, but it reflects a substantive theoretical issue about what constitutes the dataset for analysis. Those who reserve the term *bridging* for definites proceed as if there is a definites-specific phenomenon to be studied independently of cases like (1b).

We are certainly not the first to note that indefinites as well as definites may be interpreted in relation to an existing discourse entity. These observations go back at least to Hawkins (1978). In more recent work, including Asher and Lascarides (1998), Frazier (2006), and Kehler (2015), the term *bridging* is extended to these cases too, and we follow this practice in our use of the term. However, the theoretical semantics/pragmatics literature has given rather little attention to bridging of indefinites. It remains the case that bridging is primarily studied in the context of theorizing about definites. With this focus, it is natural to ignore bridging of indefinites since these NPs contain no potentially unlicensed definite marker. This results in a theoretical picture of bridging (in our sense), in which the tendency of listeners to assign bridged readings to definites is explained in terms of the requirements of the definite, while bridged readings of indefinites are attributed to general considerations of discourse coherence or plausibility. On this view, the superficial similarity between (1a) and (1b) masks a theoretically significant underlying difference: (1a) reflects a definiteness-driven phenomenon, and (1b) reflects

[1] As defined in Clark (1975), bridging inferences are implicatures driven by the interpreter's observation that a speaker has marked some part of sentence content as Given. Hence, on Clark's definition of the term, the inferences involved in (1b) should not be labelled bridging. However, as we note below and in 2.2, the term is also used in the literature in application to any cases where the hearer is required to infer some relation between entities which is not explicitly described. In fact, the term is rarely used now as originally defined by Clark, who applied it broadly to any cases where an inference is required to support the use of an expression marked as Given, including identification of the antecedent for a standard anaphoric pronoun, as well as for the Given information of an *it*-cleft.

[2] This was the position urged by one of the anonymous reviewers of this paper.

inferences motivated by general conversational considerations. On this view, as the two cases are instances of distinct phenomena, they should be studied separately. We'll call this differentiated way of thinking about bridging of definites and indefinites the *two-phenomena view* of bridging.

A different theoretical perspective emerges, though, for theorists whose starting point is the phenomenon of bridging itself, understood, as here, as the inference of specific relations between entities in discourse. As noted, rather few theorists do take this as a starting point. A notable exception is Asher and Lascarides (1998).[3] They set out to identify the circumstances in which bridging inferences are made and to provide a formal model of the inference process. In their account (to be discussed in more detail in 2.2), bridging is driven by the need to establish coherence relations between segments of text, while also satisfying other constraints imposed by lexical and compositional semantics. They argue that "it won't be possible to model all cases of bridging [in terms of] presupposition satisfaction, because bridging occurs in the absence of presupposition triggers" (p. 84). They seek to provide a unified account of bridging, an account which explains and predicts bridging of both definite and indefinite NPs. Broadly speaking, they adopt a *one-phenomenon view* of bridging, in the sense that they presume that all cases of bridging are amenable to the same *explanation*, although, of course, the details of specific cases will differ.

What, then, is the right way for theorists interested in bridging, broadly construed, to proceed? Will a more adequate model of bridging emerge by studying bridging of definites and indefinites separately? Or is an approach which treats both cases as instances of a single phenomenon more likely to be successful? And how can we tell? A significant complication in distinguishing the two approaches is that it is broadly recognized that both must account for effects of definiteness *and* various pragmatic effects, including the plausibility of postulated entity relations. Asher and Lascarides, despite their insistence that bridging cannot globally be modeled in terms of presupposition satisfaction, devote most of their paper to explaining how their framework allows for simultaneous satisfaction of coherence requirements and the presuppositions of definites. Similarly, definiteness theorists interested in bridging as a means for presupposition satisfaction recognize that it must also be constrained by expectations of discourse coherence and considerations of plausibility (see, e.g., Clark, 1975; Prince, 1992). On any view of bridging, there is, thus, some entanglement of purely linguistic and more broadly pragmatic considerations.

The distinction between the two approaches is, therefore, subtle, but seems to hinge on what is seen as the primary driver of bridging. On what we characterize as the two-phenomena view, it is linguistic form, specifically, the presence of definiteness marking, that drives bridging of

[3] They state: "We take bridging to be an inference that two objects or events that are introduced in a text are related in a particular way that isn't explicitly stated, and yet the relation is an essential part of the content of the text" (Asher & Lascarides, 1998, p. 82).

definites, while plausibility and coherence considerations drive bridging of indefinites. Under this view, definiteness is predicted to be a highly salient indicator of bridging, independently of other factors. On the one-phenomenon view, on the other hand, establishing plausible links in discourse, something required in the interpretation of any NP, is the primary driver of bridging. Definiteness marking, although it may function as a cue, is just one among several cues and is not expected to be a better indicator of an intended bridged reading than factors such as coherence and relevance. The goal of the experiments presented here is, thus, to explore whether definiteness has an independent effect on bridging. The presence of such an effect would support the two-phenomena approach. Alternatively, if there is no independent effect of definiteness, then there is no reason to treat bridging of definites and indefinites as distinct phenomena.

In order to explore this question, though, we first need a methodology to diagnose the presence of bridged interpretations. As we explain further in Section 3, prior experimental work on bridging has simply assumed that bridging arises in examples like our (1a), rather than attempting to demonstrate it. Our Experiment 1 offers a new dialogue-continuation task that provides an objective offline measure whereby bridged readings can be distinguished from non-bridged readings. Using this methodology, we probe for an independent effect of definiteness. We find no such effect. Rather, definite and indefinite NPs are bridged at equal rates, with bridging sensitive only to the presence or absence of a highly related noun (entity) in the prior discourse. This finding supports the one-phenomenon view. In Experiment 2, we turn to an online measure, assessing reading times in a novel “bridge breaking” paradigm. We, again, find no independent effect of definiteness, but we do find an interaction between definiteness and coherence: our results suggest that interpreters opt for a bridged reading most readily in the presence of two factors, a trigger NP which is definite and a context which contains a highly related anchor for bridging. This finding, which suggests a role for definiteness in either timing or strength of bridging, does not clearly distinguish the two views. We take the result to invite a (formal) model of bridging in which interpreters are responsive to multiple cues, including definiteness, which interact to determine a final interpretation.

2. Theoretical approaches to bridging and their methodological implications

In this section, we provide some more detail on the theories of definites that underlie the two approaches to bridging. The two-phenomena view derives from accounts that explain bridging of definites in terms of the general properties of definites, so to explain this view, we will have to begin by introducing some central ideas from the literature on definiteness. We do this in 2.1., where we also explain how, according to several theorists, the properties of definites are thought to trigger bridging. In 2.2., we discuss alternative views of bridging which are applicable to both definites and indefinites. We focus here on the coherence-driven account given by Asher and Lascarides (1998), which provides an explicit one-phenomenon picture of bridging.

2.1 Uniqueness and familiarity requirements as triggers of bridging

There are, broadly speaking, two families of theories of definites: uniqueness theories and familiarity theories. Uniqueness theories hold that definite NPs[4] imply, or require for their felicity, that there is a unique (salient) satisfier of the descriptive content of the definite in some contextually relevant situation. (Henceforth, we'll say that the intended referent is *situationally unique*.) The uniqueness analysis is extended to plural definites, too. In the case of plurals, the unique satisfier must be a collective entity, the mereological sum of all entities satisfying the descriptive content (Link, 1983; Sharvy, 1980). For example, sentence (3) below is predicted to be felicitous just in case there is a unique (salient) maximal set of doors in the relevant situation.

(3) The doors are open.

Moving to the second family of views, familiarity theories of definites characterize definites as carrying a familiarity presupposition or implication (Haviland & Clark, 1974), or as anaphoric (Heim, 1982; Roberts, 2003). Distinguishing uniqueness from familiarity is not always straightforward, as contextual uniqueness is often taken to be adequate to establish familiarity (see Coppock, 2022). Going in the other direction, Roberts (2003) attempts to derive the observations supporting uniqueness claims from the familiarity requirement. In contrast, Schwarz (2009) argues that definites are ambiguous between a familiarity-requiring meaning and a uniqueness-requiring meaning.

Both familiarity and uniqueness perspectives on definites have been used to argue for definiteness as a driver of bridging. Clark (1975) was first to discuss bridging, and did so in the context of his work on the Given-New Contract (Clark, 1973; Clark & Haviland, 1977). Clark is a familiarity theorist. He argues that definites signal that the intended referent is Given, which, for him, means that it can be identified with something already in the memory of the addressee, or something which the addressee can compute on the basis of their existing information. On this view, definites which lack an antecedent, such as *the door* in (1a), present a problem for the addressee: the speaker has referred to some door with a definite, signaling that it is Given, but the addressee has no relevant door stored in memory. The addressee solves this problem by making the bridging inference that the house mentioned has a door, and that *that* door is the intended referent of the definite. (Cf. Clark's (1975) discussion of "indirect reference by association".)

It should be noted from the outset that this account contains the implicit assumption that hearers prefer to resolve the apparent violation of constraints on definites by bridging, rather than by simply assuming that the referent is some door or other that they are supposed to already know about. Hearers sometimes do, in fact, interpret novel definites this way; the occurrence of

[4] We limit our discussion here to full NPs with the definite article, excluding, e.g., pronouns and demonstratives.

the house in the first sentence of (1a/b) is of this sort, and presents no problem for interpretation (at least in the current context). The hearer understands that they are supposed to envision a situation containing a unique house. This kind of solution to the problem is called *accommodation* (Lewis, 1979, a.o.). If simple accommodation were a default or preferred strategy to deal with apparent violations of definiteness constraints, with bridging as a less-preferred option, then bridging would be expected to occur more rarely, and when it does occur, some explanation would be needed for why the interpreter has deviated from the preferred strategy. So we can assume that those who propose definiteness as a trigger of bridging assume that bridging is a kind of default solution, or at least preferred over simple accommodation, given an available anchor.[5] The idea that bridging is preferred over simple accommodation is, indeed, argued for by Heim (1982, pp. 239–240), for whom this presumed preference serves an explanatory role. She considers examples such as the following:

(4) John read [a book about Schubert]$_i$ and wrote to the author$_j$.

Heim observes that this sentence can only naturally be interpreted as 'John wrote to the author of the book that he read,' and not simply as 'John wrote to some author or other that is already familiar'. She points out that this demonstrates that "accommodation in response to definites is not normally a matter of just adding the minimal amount of information that would restore felicity", but rather that a definite lacking a contextual antecedent "has to be linked by crossreferences to some already-present [discourse referent]" (p. 240).[6] Heim states the case a little too strongly here, as clearly there are interpretable cases of novel definites where bridging is not possible, as in (2) above; but Heim's example (4) is equally well explained by assuming that bridging is the default solution to an otherwise infelicitous definite, with accommodation of an unlinked discourse referent a kind of last resort.[7]

We follow the literature in assuming that whatever strategy is presumed to be involved in rendering the definite felicitous in its context, the semantics of the definite remains the same. In the theories under discussion, accommodation and bridging are interpretational strategies

[5] Strictly speaking, the inference that the house has a door is also a kind of accommodation, so, more precisely, the assumption is that accommodation of a relation between entities is preferred to what we are here calling *simple accommodation*.

[6] Heim continues: "Hence the term 'bridging': the crossreferences form a 'bridge' that connects the new discourse referent to the network of discourse referents that is already established" (p. 240).

[7] A similar idea arises in van der Sandt's (1992) account of presupposition projection and the binding of anaphora generally. Van der Sandt posits a so-called "projection line," a sequence of hierarchically ordered Discourse Representation Structures (DRSs) from the insertion point of the anaphor/presupposition to the main DRS. The system searches sequentially up the projection line for a suitable antecedent for the anaphor/presupposition; if none is found, then accommodation takes place at the highest level. This predicts a preference for binding of anaphors/presuppositions over accommodation, given an available antecedent.

triggered by the proposed semantics. The theoretical literature offers purely formal models of these different options. Following antecedents in the psycholinguistic literature, we presume these interpretative strategies to have cognitive correlates, at least with respect to mental representations of discourse content.

We have seen, then, that familiarity theorists see bridged interpretations of definites as arising as a solution to an apparent violation of the familiarity requirement. Bridging of indefinite NPs, which lack any familiarity requirement (and, indeed, are sometimes taken as markers of novelty), obviously cannot be explained in the same terms. Hence, on this view, the occurrence of bridging-like interpretations of indefinites requires a different explanation than bridging of definites. It cannot be taken to be triggered by definiteness. We emphasize that what is at issue here is what *triggers* the bridged reading, and not the (cognitive) procedure involved in establishing a bridged reading, which could be identical in both cases.

Like familiarity theorists, uniqueness theorists — those who posit that definites require situational uniqueness or unique identifiability of the referent — have also offered accounts of definiteness as a trigger of bridging. While Roberts (2003) is a familiarity theorist, she discusses the example in (5) in terms of situational uniqueness (which, in her account, is required to meet the retrievability condition on an antecedent). Imagine the following, uttered in a context where no dashboard has previously been mentioned:

(5) This car has a statue on *the dashboard*.

She points out that hearers "generally know that … there is more than one dashboard in the world," so in the absence of any situationally unique dashboard, the definite is potentially infelicitous. However, as hearers also know that there is generally only one dashboard per car, if the NP is interpreted (roughly) as "dashboard of the just-mentioned car," uniqueness is satisfied, and the felicity of the NP is guaranteed. Specifically, Roberts proposes that bridging involves reinterpreting *dashboard* as a relational noun, with (the discourse referent of) the just-mentioned car providing the second relatum.[8] For plurals, a parallel argument can be given. Consider:

(6) This car has paint on *the tires*.

Utterance of this sentence requires that there be a unique salient set of tires in the context. Here, reasoning might go as follows: there is potentially a unique maximal set of tires, namely, the set of all tires in the world, but the speaker could not plausibly intend to refer to these. The context does not make salient any other set of tires, but if *tires* is reinterpreted as "tires of the

[8] As a reviewer notes, some theorists understand the term *bridging* as describing exactly this reinterpretation strategy. This semantic analysis is equally applicable to bridging of definites and indefinites, and hence could be adopted by those favoring a one-phenomenon view. Stipulating that bridging involves nominal reinterpretation, therefore, does not decide the question of whether bridging of definites and indefinites is *triggered* in the same way or not.

just-mentioned car", then once again, a single set of tires becomes salient, namely, those of the car. Hence, *tires* is given a bridged interpretation, satisfying (plural) uniqueness.

A similar account is given in Kehler (2015) of the example below:

(7) The politician approached the courthouse and proceeded up the steps.

Kehler explains the bridging as follows: "On a plausible analysis of such cases, the hearer accommodates (Lewis, 1979) the existence of the referent… The hearer, confronted with a referring expression that requires a uniquely identifiable referent, infers the existence of courthouse steps to meet that constraint." (Kehler, 2015, p. 634).

Thus, both familiarity theories and uniqueness theories of definites provide explanations of the discourse conditions under which bridging will occur, namely, in order to satisfy the otherwise unmet felicity conditions on the use of a definite. In both cases, the explanation does not extend to bridging of indefinites, so despite the surface similarity of bridging of definites and indefinites, different explanations are required.

2.2 Extending bridging to indefinites: The one-phenomenon view of bridging

The observation that indefinites can bridge (in our sense) goes back at least to Hawkins (1978, pp. 173ff.), who treats them as parallel to the associative anaphoric uses of definites. Hawkins' examples include these:

(8) Fred bought a book from Heffer's. He was dismayed to find that a page was torn/some pages were torn.

(9) I've just inspected a house. I decided not to buy it because a window was loose/some windows were loose.

Hawkins proposes that associative links involving indefinites are subject to the same (complex and difficult to specify) pragmatic constraints that govern associative anaphora with definites, but observes a contrast in the relation to uniqueness: the felicitous use of an associative indefinite requires that the listener can assume that the intended referent (singular or plural) is *not* situationally unique. Hence, (8) is felicitous with the singular *a page*, because books typically have more than one page, and also with the plural *some pages*, as books typically have more than two pages (hence, the speaker can use *some pages* to describe a plural subset of the total set of pages).

Hawkins observes that indefinites with a preceding "trigger" (in our terminology, an *anchor*) can often be understood either as associated with the trigger or not. (See also the discussion in Kehler (2015, p. 635) regarding his example (34).) In the following example, Hawkins claims, it's unclear whether the tires that Fred sold came from the car he bought, or not:

(10) Fred bought a car last week, and then he sold some tires to his friend.

Hawkins holds, though, that there are cases where an indefinite is obligatorily interpreted associatively, namely, where the context satisfies the non-uniqueness requirements that he assumes and where the "pragmatics of the remainder of the sentence" force this reading, as intuitively is the case in (8) and (9) above. Hawkins' discussion, thus, suggests a commonality between the bridged, or associative, readings of definites and indefinites. Both stem, somehow, from the associations arising from the mention of the anchor noun.

To our knowledge, the only established formal model of bridging which recognizes a commonality between bridging of definite and indefinite NPs is that of Asher and Lascarides (1998). Asher and Lascarides argue that bridging is a consequence of, and subserves, the construction of coherent discourse structure, an idea which originates with Hobbs (1979). On their view, constructing discourse structure requires the establishment of rhetorical relations between each Elementary Discourse Unit (typically, but not always, a clause; see Asher & Lascarides, 2003). They propose that bridging relations are constructed where they are necessary to support a plausible coherence relation or, alternatively, where the relation between the entities is implied by the most plausible coherence relation available. Let's illustrate this with the basic examples from the beginning of the article, repeated here:

(1) a. Yasmin approached the house. The door was open.
b. Yasmin approached the house. A door was open.

On the Asher and Lascarides account, the interpreter of the sequences in (1) is in search of a plausible rhetorical relation between the two segments. For both examples, the most plausible relation would be Elaboration or Background. But in order for the second segment to elaborate on the eventuality described in the first, the door under discussion must be a door related to the eventuality of Yasmin approaching the house. Given world knowledge that houses typically have doors, it is straightforward for the listener to infer that the door in question belongs to the house introduced in the first segment. As houses can have a unique (or uniquely salient) external door, use of the definite is felicitous and plausibly leads the hearer to further infer unique instantiation. As houses also can have multiple external doors, use of the indefinite (which can, in turn, be seen as avoidance of the more restricted definite) implies that uniqueness is not satisfied, so the hearer may (but need not) envision multiple doors in the situation. Thus, on the Asher and Lascarides view, definiteness has a clear role to play in interpretation overall, but is not the feature responsible for triggering bridging. Indeed, we should emphasize here that rejection of the claim that definiteness is the primary trigger of bridging does not entail a rejection of the claims surveyed in the previous section about the function of, or constraints imposed by, definites (or indefinites).

The approach advocated by Asher and Lascarides, as well as Hawkins' earlier discussion, demonstrates that a theory can fully represent the role of definiteness (and indefiniteness) in supporting and constraining bridging without taking definiteness to be the trigger of bridging. Suppose that interpreters encountering a novel noun with a plausible anchor already in the discourse typically build the relation to the anchor into their initial interpretation, irrespective of definiteness. But interpreters also seek to satisfy any requirements created by the use of a definite (or indefinite), such as a requirement for situational uniqueness (or non-uniqueness). On this account, bridging driven by an observation of plausible relations between entities *facilitates* satisfaction of the requirements of a definite, by making salient a uniqueness-satisfying relation. Thus, definiteness may reinforce or strengthen the likelihood of a bridged reading without being its immediate trigger.

This kind of view is consistent with conclusions drawn by Schumacher (2009) from an ERP study of the interpretation of definites and indefinites. In her study, she investigated the ERP signatures associated with the interpretation of definite and indefinite NPs in three conditions: repetition of a previously used noun (Given), a new noun highly related to prior context (Inferred), and an unrelated noun (New). Schumacher explored both early (N400) and later (P600) ERP effects. With respect to the N400, Schumacher found no significant effect of definiteness; on this basis, she argues that in early interpretation, hearers attempt to integrate new NPs into existing representations regardless of definiteness. It is only with respect to the P600 that any effect of definiteness occurs. Schumacher argues that "the linking attempts occur for both definite and indefinite entities and are governed by the fit of the head noun with the information provided by the context. This suggests that the underlying processes are guided by coherence constraints and that a strict correspondence between definiteness marking and integration processes cannot be maintained" (p. 100).[9] Thus, Schumacher's data suggest that the search for an antecedent in context is not motivated by the presence of a definite. Consequently, her conclusions can be read as supporting a one-phenomenon view of bridging, a view according to which the bridging of both definites and indefinites is triggered in the same way – by general expectations of, or preferences for, links between prior discourse and new information.

The experiments we report here similarly explore the role of definiteness in bridging, but using a different set of methodologies that explore later stages in interpretation. A key observation underlying our experimental methodology is that there are cases where definiteness and more general considerations of coherence push in different directions. These cases are particularly useful for exploring whether definiteness plays an independent role in triggering bridging. Consider, then, the examples in (11), where the subject NPs *the/a doctor* are not highly related to any entities mentioned in the context given:

[9] Schumacher's analysis reflects her interpretation of the N400 as a marker of integration. We note that other interpretations are offered in the literature (see the review in Li & Ettinger, 2022).

(11) a. Justin walked down to the stable. The doctor was there.
b. Justin walked down to the stable. A doctor was there.

Regardless of the identity of the doctor, the second sentence here is interpretable as cohering with the first via the Elaboration relation. Establishing this relation doesn't require assuming any particular relation between the doctor and anything mentioned in the first sentence. As far as coherence establishment is concerned, there is no more reason to posit a bridge in the definite case than in the indefinite case. With respect to interpretation of the definite, bridging is not *obligatory* for felicity; the interpreter can simply assume that the definite *the doctor* refers to some assumed-familiar referent, as with *the stable* in the initial sentence. But, as also noted earlier, theories of bridging as definiteness-driven seem predicated on the assumption that bridging (accommodation of a between-entity relation) is preferred over simple accommodation of a new referent as a strategy for meeting the requirements of definites. The two-phenomena view of bridging would, therefore, seem to predict an increased tendency to attempt to bridge *doctor* to the content of the previous sentence in the definite case, over the indefinite case. It is this prediction that will be tested in the experiment described in Section 4, where we test for an independent effect of definiteness that persists even in passages that lack a highly related entity to serve as an anchor.[10]

3. A lacuna in prior experimental work on bridging

To address the question of interest, it will be necessary to determine whether experiment participants assign bridged interpretations to target items. This turns out to require the development of novel methodologies, as in prior experimental work involving bridged interpretations, bridging is assumed, rather than directly diagnosed. Indeed, in prior experimental work, NP bridging has tended to be an experimental tool, rather than the focus of investigation in its own right. In a very early, much cited study, Haviland and Clark (1974) used definite NPs (among other presupposition triggers) to test their Given-New theory of sentence comprehension. Haviland and Clark presented participants with what they called *Direct Antecedent* stimuli, like (12a), or *Indirect Antecedent* stimuli, like (12b) (cf. Schumacher's Given and Inferred conditions).

(12) a. We got some beer out of the trunk. The beer was warm.
b. We checked on the picnic supplies. The beer was warm.

[10] One of our anonymous reviewers remains unconvinced that the discussion of accommodation versus bridging in the theoretical literature entails an 'always try bridging first' position. The reviewer's position is that this literature entails only that bridging is preferred (or perhaps obligatory) if there is a "natural bridge" available. However, from a processing perspective, we assume that a listener attempting to ascertain whether there is a natural bridge available must *attempt* bridging, meaning that a definite would, indeed, always trigger a bridging attempt. But we acknowledge that this is an untested assumption. Readers who share the reviewer's skepticism about our interpretation of the theoretical literature may simply consider the experiments presented here as seeking to distinguish two alternative views of the processing of bridging.

Participants were asked to respond when they felt sure that they understood the second sentence. Participants responded more quickly in the Direct Antecedent case than the Indirect Antecedent case. Haviland and Clark take this as evidence that in the Indirect Antecedent case, participants are engaging in "bridge building" (inferring, e.g., that the picnic supplies included beer), accounting for the extra processing time. (An effect of mere repetition was ruled out by a separate experiment.) Interestingly, Haviland and Clark don't consider the possibility that the extra processing time might have arisen from participants simply accommodating a referent without bridging (just as they would need to do to interpret the definites in the first sentence of the prompt). Accommodation of a referent could also lead to an increase in processing time. It certainly seems unlikely that participants would fail to bridge; however, the experiment does not offer any *direct* evidence of the interpretations constructed.[11]

This turns out to be a common feature of much experimental work on NP interpretation in the decades that have followed, even though many more sophisticated measures have been developed and deployed. For example, Clifton (2012) reports a study (utilizing both reading time and eye movements) intended to investigate the effects on interpretation of the purported uniqueness/non-uniqueness implications of definite and indefinite determiners. Clifton presented participants with sets of examples such as the following:

(13) In the kitchen/the appliance store, Jason checked out a/the stove very carefully.

These stimuli begin with a context phrase which includes a potential bridging anchor for the target NP introduced later in the sentence. The target is always highly related to the potential anchor (both kitchens and appliance stores typically contain stoves), but the contexts differ as to whether the target noun would typically be uniquely instantiated in that context or multiply instantiated. Clifton assumes that the relatedness between the noun in the context phrase and the target noun will lead to a bridging attempt, regardless of the definiteness of the target NP; articles are expected to affect bridging only to the extent that their implications of (non)-uniqueness may fail to match those of the context.[12] (Again, this is in line with Schumacher's ERP findings.) Clifton found slower reading times in the mismatched conditions in some versions of the experiment, but, as in the earlier Haviland and Clark study, did not attempt to ascertain what interpretations participants actually arrive at for the target NPs. Clifton, indeed, notes that his data "cannot unambiguously determine whether the slowed reading reflected time taken to

[11] An anonymous reviewer suggests that the fact that the definite *the beer* is felicitous, despite the lack of an overt antecedent, "establishes the presence of bridging", which, they suggest, would explain why Haviland and Clark do not consider the option of participants failing to bridge. However, as discussed above, accommodation of a referent without a relation to an existing discourse referent is always an option. In both example stimuli, the first sentence contains an antecedentless definite, which on the Clarkian view could only be interpreted via accommodation.

[12] Thus, Clifton is another example of a researcher who uses the term *bridging* in the broader sense that we use it, and, in fact, this usage seems to be common in the psycholinguistic literature.

accommodate the presuppositions or simply disruption triggered by noting that presuppositions had not been met" (p. 497). So, although the experiment is designed to encourage a bridged interpretation of target NPs, it does not provide any direct evidence of the actual interpretations generated by participants.

We address this methodological lacuna in the experiments described in the next sections.

4. Experiment 1: "Which one?" dialogue continuations

This experiment aims to address the "one-phenomenon-or-two" question by exploring whether an independent effect of definiteness can be observed in the rate at which participants assign bridged interpretations in stimuli in which definiteness and availability of a high-related anchor are manipulated. Recall, from the end of 2.2, the observation that although the one-phenomenon and two-phenomena views of bridging both predict effects of contextual plausibility and definiteness, the two differ with respect to predictions involving definites in contexts where bridging does not straightforwardly support coherence. The two-phenomena view, according to which definiteness is a principal driver of bridging, suggests that when an interpreter encounters a novel definite, they will attempt to bridge, and will be more likely to adopt interpretations requiring bridging to a less-plausible anchor than in the case of a novel indefinite. Consequently, in a comparison of definites and indefinites in the absence of a highly related potential anchor, we would expect higher rates of bridging for definites than for indefinites. We test this prediction in this experiment.

The approach requires us to determine whether a participant has assigned a bridged reading to a given NP, which, in turn, requires the development of new methodology. The methodology was developed with two desiderata in mind. First, we wanted a measure of bridging that did not involve experimenter interpretations of participant responses. Second, we wanted to avoid prejudging what bridged readings were possible for a given stimulus, as different hearers may infer different relations. For example, given (14), one interpreter might understand *the colors* to be the colors of the leaves, while another might infer that Liping is looking at a broad outdoor scene and that *the colors* refers to the colors throughout the scene.

(14) Liping watched the leaves falling from the trees. The colors were beautiful.

While the former reading seems more plausible, we wanted to ensure that our measure of bridging did not prejudge which inferred relations would count as bridging.

To accomplish this, we introduce a new task in which participants are asked to read a passage that contains a potentially bridged NP (the *target*) and to answer a question that probes how the participant, in fact, interpreted it. In this task, the question is embedded in an exchange between two speakers, as in (15).

(15) Speaker A: Nigel and I went out last night to that new restaurant. The waiter was very friendly.
Speaker B: Wait, sorry, I wasn't listening. Which waiter are you talking about?
Speaker A: ___

The strategy used here is indirect but, we believe, can reliably reveal the participant's interpretation of the target NP. We presume that to answer the clarification question on behalf of Speaker A, the participant will draw on whatever they take Speaker A to have meant, which is precisely the information we are after. This indirect strategy is preferable to a direct inquiry to participants about their interpretation, which would likely lead to overthinking, and also would quickly reveal the experimental goal. Additionally, while being asked to explain or describe an interpretation is not a normal discourse task, responding to a clarification question is, which we take to be an additional benefit of the strategy.

To assess bridging, we measure the degree to which participants use material from the context sentence in responding to the "which N?" prompt. Specifically, we use a measure of string similarity to test the amount of string overlap between the participant responses elicited by passages like (15) and the context sentence of the prompt (*Nigel and I went out last night to that new restaurant*). We assume that a strong context~response similarity reflects (a report of) a bridged interpretation. We do not posit a numerical boundary or cut-off for bridging. The response *the waiter at the restaurant where Nigel and I went out last night* will produce a higher context~response similarity score than responses like *the one who I saw* or *the friendly one.* The first response more clearly indicates that the participant has interpreted the new entity in relation to entities previously mentioned than do the others. We use passages constructed in a prior elicitation task (see 4.2 below), where naive participants generated continuations for a set of prompts. These continuations provided the stimuli for the current experiment, in which we manipulate the definiteness of the target noun and the presence or absence of a highly related noun in the context sentence.

Under the one-phenomenon view of bridging, the key driver for a bridged interpretation is entity relatedness, not definiteness, whereas the two-phenomena view attributes bridging of definites to the presence of a definite. Of particular interest is the comparison between definite and indefinite NPs in the low-related condition, where the two views make competing predictions, as discussed above. Experimentally, the two-phenomena view predicts an independent effect of definiteness, while the one-phenomenon view does not.

4.1 Participants

For the main task, seventy-two participants who had IP addresses in the United States or Canada were recruited via Amazon Mechanical Turk. All participants were paid $10 for a task estimated to take under an hour. We only used data from participants who answered "no" to a question in a

background questionnaire that asked whether any language besides English was spoken at home before the age of 6, in order to restrict ourselves to analysis of data from monolingual native-speaker participants. The data for the analysis comes from 55 monolingual English-speaking participants.

4.2 Materials

The materials followed the dialogue structure shown in (15), with the content of the first turn determined from an elicitation task with a separate set of participants.

For the elicitation task, we recruited a separate set of 72 English-speaking participants, none of whom participated in the main task and all of whom were paid $13 for a task estimated to take an hour. We used the data from 54 monolingual English-speaking participants. Participants accessed a web-based interface where they wrote story continuations for prompts like those in **Table 1**.

Table 1: Experiment 1 elicitation task materials consisting of a context sentence and a prompt NP (the subject of the next sentence, which participants were to complete).

High-related/definite: Nigel and I went out last night to that new restaurant. The waiter …
High-related/indefinite: Nigel and I went out last night to that new restaurant. A waiter …
Low-related/definite: I was shocked by something I saw yesterday on the news. The waiter…
Low-related/indefinite: I was shocked by something I saw yesterday on the news. A waiter…

For each item set in the elicitation task, we manipulated Definiteness and Relatedness: each of the 40 target nouns (*waiter* in **Table 1**) appeared in either a definite or indefinite NP and was preceded by a context sentence that mentioned either a high-related anchor (*restaurant*) or contained only low-related nouns (here, *news*). The materials were constructed so that each noun and each context sentence appeared in all four conditions. In other words, not only did the target noun *waiter* appear in all 4 conditions, as shown in **Table 1**, but the high-related context sentence in **Table 1** also appeared as a low-related context sentence for another noun (*Nigel and I went out last night to that new restaurant. The report…*), and vice versa for the other context sentence (*I was really shocked by something I saw yesterday on the news. The report…*). This counterbalancing ensures that in the main task, any evidence for bridging in the high-related condition could not be attributed to an independent effect from the context sentence. For example, it is possible that some context sentences have content that is more likely to be re-used

in a dialogue continuation; maybe going out to a restaurant is inherently more interesting to talk about than seeing something in the news, or vice versa. See Appendix A for further information about the elicitation task, with the full set of elicitation materials available in the Supplementary Materials.[13]

To construct the main task materials, we sampled the set of continuations produced by participants in the elicitation task. All 40 target NPs from the elicitation task were represented, and we selected 4 continuations written in the high-related condition and 4 written in the low-related condition. These included both singular and plural target nouns (26 singular, 14 plural). We independently manipulated the definiteness of the target NP to yield a set like those shown in **Tables 2** and **3** for each target NP.

Table 2: Experiment 1 main task materials for a singular target NP, consisting of Speaker A's turn (context sentence + one of four continuation sentences), Speaker B's question ("Wait, sorry, I wasn't listening. Which X are you talking about?"), and a prompt for participants to fill in Speaker A's reply.

High-related condition:
Speaker A: Nigel and I went out last night to that new restaurant. 1. [The/A] waiter was very friendly. 2. [The/A] waiter welcomed us and took our order. 3. [The/A] waiter was friendly and helpful. 4. [The/A] waiter recommended the best entree. Speaker B: Wait, sorry, I wasn't listening. Which waiter are you talking about? Speaker A: _____
Low-related condition:
Speaker A: I was really shocked by something I saw yesterday on the news. 1. [The/A] waiter suddenly went on a racist rant and it was all caught on video. 2. [The/A] waiter punched the customer in the face. 3. [The/A] waiter poisoned the food. 4. [The/A] waiter stole a customer's credit card. Speaker B: Wait, sorry, I wasn't listening. Which waiter are you talking about? Speaker A: _____

[13] The following OSF link contains the materials and instructions for the elicitation task, as well as full materials, data, R scripts, and post-hoc analyses for all other experiments reported in this article: https://osf.io/6w7v8/?view_only=dbbccc0a5d83482eb998bd17e539b409.

Table 3: Experiment 1 main task materials for a plural target NP, following the format for Table 2.

High-related condition:
Speaker A: Barbara was grinning from ear to ear when she walked into her classroom. 1. [The/Some] desks were arranged in the shape of a heart. 2. [The/Some] desks were finally arranged in a circle the way she wanted. 3. [The/Some] desks were set up backwards. 4. [The/Some] desks had moved around. Speaker B: Wait, sorry, I wasn't listening. Which desks are you talking about? Speaker A: _____
Low-related condition:
Speaker A: Charles has a large table top aquarium. 1. [The/Some] desks were neatly organized and carefully placed so that the aquarium wouldn't knock over. 2. [The/Some] desks are pushed together to make a stand for it. 3. [The/Some] desks were pushed together in order to accommodate it. 4. [The/Some] desks were pushed together to accommodate the large tank. Speaker B: Wait, sorry, I wasn't listening. Which desks are you talking about? Speaker A: _____

Using elicited continuations, as we did here, avoids experimenter bias in the readability of the high- versus low-related conditions. The manipulation of definiteness means that roughly half the time, the Experiment 1 passage deviates from the elicitation participant's original version (as elicitation participants were responding either to a definite or an indefinite prompt). However, we note that there were over a dozen cases where the elicitation data included the same continuation in both the indefinite and definite conditions, suggesting that these continuations are amenable to a definiteness manipulation (e.g., *Hilda created a nice arrangement of fruit. [The,A] banana was the centerpiece*). We included a number of those cases in the Experiment 1 materials, yielding a distribution across the continuations selected from the elicitation task such that half the continuations were from underlying indefinite prompts, and half were from underlyingly definite prompts, with 5 produced for both indefinite and definite prompts.

For the choice of continuations, we randomly selected NP continuations from the elicitation dataset and then eliminated selections and made new selections according to the following criteria: no continuations with modifiers (e.g., *The waiter who brought our food was rude*), no continuations with non-referential NPs (e.g., *A rug would make the room look better*), no continuations with additional common ground assumptions (e.g., *beforehand* in *The lion had attacked some people beforehand*), no continuations with idiomatic expressions that require a

definite (e.g., *The clouds are few and far between*), no continuations that sounded awkward or confusing to a native speaker (e.g., *The page was given in the mail for coupon deals at the grocery store*), and no continuations for which the definiteness manipulation produced incoherence (e.g., *Some leaves are brown and some are green*, which becomes contradictory if the first NP is definite as in *The leaves are brown and some are green*). Replacements were selected to balance the overall number of continuations that had been originally definite/indefinite in the elicitation task; the materials for any given target NP included at least one continuation that had been elicited from a definite prompt, and at least one from an indefinite prompt. We fixed small typos (e.g., *it's* → *its*, *wont* → *won't*). The only other alteration we made to the elicited continuations was to use *the/some* as the definite/indefinite alternation for NPs with the target noun *fish* (instead of *the/a*); this was because almost all of the continuations with this target noun with the definite in the elicitation task had been assigned a plural interpretation. (See the Supplementary Materials for the full set of materials).

Eight lists were constructed, such that participants saw each of the 40 target NPs once in a single passage (e.g., *waiter* in one of the 8 variants in **Table 2**). Each list included a particular context sentence only once (i.e., a participant who saw *Nigel and I went out last night to that new restaurant* with the target NP *the waiter* did not see that context sentence with its low-related NP *the report*). Each list contained an equal number of high-related and low-related items and an equal number of definites and indefinites. In addition to the 40 target items, each list contained 40 filler passages which were similar to the target dialogues: Speaker A's first turn consisted of two sentences; Speaker B asked a clarification question (e.g., *What did you say? Why is school closed?*, or *Sorry, someone was talking to me. What was on the test?*), and the participant was asked to fill in Speaker A's reply.

4.3 Procedure

Participants accessed the experiment via a web-based interface linked from the Mechanical Turk environment. For all experiments reported here, participants gave informed consent before proceeding to the task. Each item was displayed on a separate page. Participants were instructed to fill in the final utterance in the dialogue. (See the Supplementary Materials for full detail.)

4.4 Results

For the 2200 responses collected, we used a string similarity metric as a proxy for participants' assignment of a bridged interpretation. The participant responses in (16) and (17), shown underlined, illustrate the contrast we are aiming to capture. In (16), the participant is clearly linking the newly mentioned waiter to the context provided by the initial sentence; this is shown by the re-use in the response of multiple words from the context sentence. In (17), in contrast, the participant does not re-use words or content from the context sentence, but instead gives a rather

under-informative response, simply repeating the content that is asserted in the continuation sentence. We take this to indicate that the participant has assigned a non-bridged interpretation to the NP – plausibly, the participant doesn't have any clear idea of what lion is being talked about.

(16) Speaker A: Nigel and I went out last night to that new restaurant. A waiter was friendly and helpful.
Speaker B: Wait, sorry, I wasn't listening. Which waiter are you talking about?
Speaker A: A waiter at that new restaurant I went to last night with Nigel

(17) Speaker A: There's been a lot of controversy recently at my university. The lion was used as a mascot.
Speaker B: Wait, sorry, I wasn't listening. Which lion are you talking about?
Speaker A: The lion used as a mascot

For the string overlap measure, we compute a score for the context~response sentence overlap as a proxy for bridging. We also compute the overlap between the participant's response and the continuation sentence, with the hypothesis that in unbridged cases like (17), the participant may resort to reasserting information from the continuation sentence in order to answer the clarification question. For any pair of sentences, the similarity measure was computed by treating each sentence as a "bag of words" represented as a vector in a multi-dimensional lexical space. We then evaluated the distance between those two vectors. Pre-processing of the sentence text was conducted using the Python Natural Language Toolkit (nltk; Bird et al., 2009). Specifically, the bag of words for a given sentence consisted of all the words in that sentence after we excluded punctuation and stop words (e.g., *and, that, was*; nltk.corpus' stopwords. words("english")).

We present two analyses – one with stemmed words (e.g., *cats/cat* both appear as *cat*; nltk PorterStemmer) and one with lemmatized words (e.g., *buy/buys/buying/bought* all appear as *buy*; nltk WordNetLemmatizer). The distance between two sentences was measured as the cosine of the angle between their respective word vectors (Manning & Schütze, 1999). With this measure, two sentences that consist of an identical set of words have value 1, and two sentences that have no overlapping words have value 0.[14] Appendix B includes a sample of participant responses and the associated similarity scores that were computed for the context~response comparison and the continuation~response comparison.

Table 4 and **Figure 1** show the mean similarity scores between the context sentence and the participant response. There are higher scores for high-related NPs than low-related NPs, and little difference by Definiteness. We construct a linear mixed-effects regression to predict the

[14] As an illustration of the cosine similarity metric, the strings in (i) and (ii) are shown with their bag-of-words representations, which yield the associated cosine similarity scores between 0 and 1:

similarity score with fixed effects of Relatedness and Definiteness and their interaction (coded as –.5/+.5 low/high and –.5/+.5 indefinite/definite). The model contained random effects for participants and NPs, with random intercepts and slopes. The converging model for stemmed similarity contained only by-participant random slopes for Relatedness and Definiteness and a by-NP random slope of Relatedness (see Barr et al., 2013, for model convergence approaches). For lemmatized similarity, the converging model contained full by-participant random-effects structure, but only a by-NP random slope of Relatedness. The results show a significant main effect of Relatedness (stemmed: $B = 0.147$, $SE = 0.02$, $t = 7.08$, $p < 0.001$; lemmatized: $B = 0.141$, $SE = 0.02$, $t = 6.688$, $p < 0.001$). There is no main effect of Definiteness nor a Relatedness × Definiteness interaction (p's > 0.3).

Table 4: Mean similarity scores between the context sentence and the participant response (higher similarity is taken to be a proxy for bridging between the target NP and the context sentence).

	Stemmed similarity		**Lemma similarity**	
	low	**high**	**low**	**high**
indefinite	0.20	0.36	0.18	0.33
definite	0.22	0.35	0.19	0.32

(i) a. Context sentence: *Nigel and I went out last night to that new restaurant.*
b. Response to 'Which waiter?': *A waiter at that new restaurant I went to last night with Nigel*
c. Similarity calculation
- Vocabulary across a&b without stop words: ['last', 'new', 'nigel', 'night', 'restaurant', 'waiter', 'went']
- vector for (a): [1, 1, 1, 1, 1, 0, 1]
- vector for (b): [1, 1, 1, 1, 1, 1, 1]
- cosine(a,b) = 0.93

(ii) a. Context sentence: *Nigel and I went out last night to that new restaurant.*
b. Response to 'Which waiter?': *The one at the new restaurant in town.*
c. Similarity calculation
- Vocabulary across a&b without stop words: ['last', 'new', 'nigel', 'night', 'one', 'restaurant', 'town', 'went']
- vector for (a): [1, 1, 1, 1, 0, 1, 0, 1]
- vector for (b): [0, 1, 0, 0, 1, 1, 1, 0]
- cosine(a,b) = 0.41

The intuition behind the cosine metric is that the sentences are treated as vectors in a multi-dimensional "vocabulary space" where each word is an axis and the angle between the two sentence vectors represents their similarity: two sentences are similar if their vectors deviate by only a small angle (large cosine value); two sentences are dissimilar if their vectors deviate by a large angle (small cosine value).

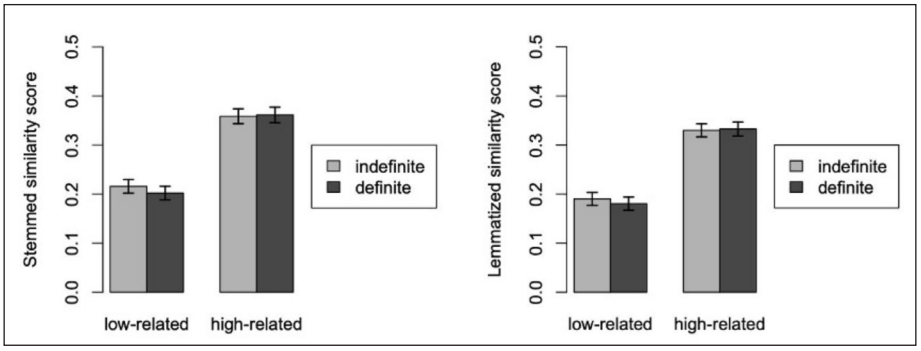

Figure 1: Context~response similarity (bridging proxy) as a function of Relatedness and Definiteness on two types of similarity scores (stemmed on left and lemmatized on right); error bars show standard error over by-participant means.

Table 5 and **Figure 2** show the similarity scores between the continuation sentence and the participant response. This measure doesn't index bridging and, instead, may capture ways in which participants seek to link an entity to the context when a bridged interpretation is not favored. The results show the highest similarity scores in the low-related indefinite condition. The converging model for stemmed similarity contained only by-participant random slopes for Relatedness and the Relatedness × Definiteness interaction and a by-NP random slope of Relatedness. For lemmatized similarity, the converging model contained only by-participant and by-NP random slopes for Relatedness. The results show a significant main effect of Relatedness, whereby the continuation~response similarity is higher for the low-related condition than the high-related condition (stemmed: $B = -0.077$, $SE = 0.016$, $t = -4.810$, $p < 0.001$; lemmatized: $B = -0.077$, $SE = 0.016$, $t = -4.942$, $p < 0.001$), and a less reliable effect of Definiteness, whereby definites have lower scores than indefinites (stemmed: $B = -0.0181$, $SE = 0.009$, $t = -2.001$, $p < 0.05$; lemmatized: $B = -0.0154$, $SE = 0.09$, $t = -1.723$, $p = 0.09$), which together are driven by a significant Relatedness × Definiteness interaction (stemmed: $B = 0.052$, $SE = 0.019$, $t = 2.693$, $p < 0.001$; lemmatized: $B = 0.054$, $SE = 0.019$, $t = 2.894$, $p < 0.005$).

Table 5: Mean similarity scores between the continuation sentence and the participant response.

	Stemmed similarity		**Lemma similarity**	
	low	**high**	**low**	**high**
indefinite	0.41	0.31	0.40	0.30
definite	0.37	0.32	0.37	0.31

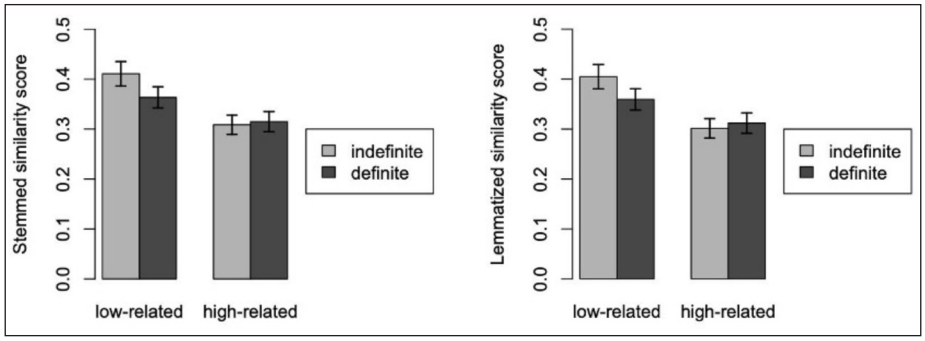

Figure 2: Continuation~response similarity as a function of Relatedness and Definiteness on two types of similarity scores (stemmed on left and lemmatized on right); error bars show standard error over by-participant means.

In reviewing our materials, it is apparent that our items are somewhat heterogeneous. We, therefore, undertook several post hoc analyses (see the Supplementary Materials) to test the stability of our findings across different subsets of the data, focusing on whether an independent effect of definiteness is apparent under a different treatment of our data. We compare singular vs. plural target nouns and relational vs. non-relational target nouns,[15] and we eliminate additional cases with non-referential nouns. Lastly, we report a uniqueness rating study (see Appendix C) that we conducted to test that our items satisfied our intended constraint that all context~noun pairings allow an interpretation in which the noun (or the entity it corresponds to) could plausibly be either uniquely or non-uniquely instantiated in the situation given by the context. The only problem stimuli we identified in the rating study were two cases where the likelihood of uniqueness was low. However, in both of these cases, the target noun is robustly relational (target noun *leg*, given *chair* in context, and target noun *page*, given *book* in context) and clearly allows for use of the definite, despite contextual non-uniqueness. We conclude that our stimuli do satisfy the intended constraint.[16] The upshot of these post-hoc analyses is that

[15] Definites with relational noun heads are argued in the literature to have some special behaviors which are relevant to bridging. Barker (2005) argues that relational nouns may head definites which *lack* a uniqueness (or familiarity) requirement, a class of cases which he dubs *weak definites*. Weak definites may be given bridged readings; as with indefinites, the explanation of these readings cannot proceed via the need to satisfy uniqueness.

Schwarz (2009) also discusses bridging of relational nouns, arguing that bridging of (in his terminology) *strong* (or *anaphoric*) *definites* is allowed only with relational nouns. Schwarz further argues that bridged anaphoric definites *do* give rise to a uniqueness requirement (in contrast with non-bridged uses of anaphoric definites). Schwarz's analysis pertains to German data. It's unclear whether the contrast between Schwarz's claims and Barker's reflects a language difference, or simply points to further complications in the data. It's also unclear how the relational/non-relational distinction would be predicted to affect bridging judgments in our data, but given that some theorists take this to be a theoretically significant distinction, we chose to carry out this post hoc analysis.

[16] Thank you to our reviewers for raising questions about the Experiment 1 and Experiment 2 materials, which led to the norming study and these additional analyses.

the findings from the main analysis are robust. The entity relatedness effect on bridging (as measured via context~response similarity) is consistently significant, and, crucially, no evidence of an independent effect of definiteness emerges.

4.5 Discussion

As we argued earlier, an independent effect of definiteness on bridging would support two-phenomena approaches, which take there to be a specific definiteness-driven phenomenon of bridging distinct from coherence-driven bridging. In contrast, the absence of such an effect supports a one-phenomenon approach like that of Asher and Lascarides (1998), which takes there to be a single, broad phenomenon which encompasses bridging of both definite and indefinite NPs. Experiment 1 shows no independent effect of definiteness, supporting the one-phenomenon view. The experiment suggests that in interpretation, recognition that a newly mentioned entity is highly related to some entity already in the discourse model will lead to a bridged interpretation of the new noun, regardless of definiteness marking. It suggests further that bridging is available as a way of satisfying the uniqueness/familiarity requirements of a definite only in the specific case where bridging is already supported by contextual factors.

Besides the findings related to context~response similarity, we saw that the continuation~response similarity scores show a marked increase for low-related NPs, and particularly for indefinites in that condition. Some cases of high continuation~response similarity may be coming from participants who have not bridged the target noun. Faced with the "Which X?" question, a non-bridging participant has two choices: to invent properties of an imagined brand-new referent, or to use material from the continuation sentence. The observation that particularly high rates of continuation~response similarity occur with low-related indefinites may reflect the fact that no feature in the context sentence inclines the interpreter towards a bridged interpretation. And, in fact, the effect of indefiniteness here is quite weak. When we compare the effect of Relatedness on bridging from the context~response similarity analysis with the Relatedness $\times$ Definiteness interaction in the continuation~response similarity analysis, the former represents a difference of at least 0.14 in the 0 to 1 scores (see **Table 4**), whereas the latter represents a difference of at most 0.04 (def vs. indef in the low-related condition, see **Table 5**). We take this to suggest that the extent to which bridged interpretations are sensitive to relatedness (with no influence from definiteness) far outweighs the extent to which definiteness influences participants' strategies in potential cases of non-bridging in the low-related condition.

Before concluding Experiment 1, we turn to a possible alternative explanation of context~response similarity. It is possible that high string overlap between the participant's response and the context sentence occurs in the high-related condition even in the absence of bridging, simply because of the high relatedness of the target noun and its potential anchor. For

example, high overlap could be driven by responses like the (non-attested) Speaker A response in (18). Here, the word *restaurant* is used, not because the context provides it as the anchor for a bridge, but because restaurants are common places to find waiters.

(18) Speaker A: Nigel and I went out last night to that new restaurant. A waiter was friendly and helpful.
Speaker B: Wait, sorry, I wasn't listening. Which waiter are you talking about?
Speaker A: A waiter at a restaurant I once went to.

In order to rule out this alternative explanation for the high similarity scores between the context sentence and the response, we conducted an additional check, with the aim of ascertaining, for example, how often participants use *restaurant* in a response about a waiter, even if the context sentence doesn't contain the word *restaurant*. For each target noun, we checked the frequency of occurrence of the high-related anchor noun (e.g., *restaurant* for *waiter*, *classroom* for *desks*, *arrangement of fruit* for *banana*, *room* for *window*, *wedding* for *guests*) in responses in both the high-related condition, in which the high-related noun occurred in the context sentence, and in the low-related condition, which did not contain the high-related noun in the context sentence. **Table 6** reports the by-condition averages of the rate of mention of the related word for each target NP.

Table 6: Proportion of mentions of the related noun in participant responses.

	low	**high**
indefinite	0.09	0.70
definite	0.08	0.68

While participants do occasionally use the related word in their responses even when it has not occurred in the context sentence (i.e., in the low-related condition, where the use rates are 0.09 and 0.08), they do so at a much higher rate when the related word is present as a candidate anchor in the context sentence (i.e., in the high-related condition, where the use rates are 0.70 and 0.68). An analysis of the binary outcome of the mention of the related word shows only a main effect of Relatedness ($B = 4.844$, $SE = 0.506$, $z = 9.581$, $p < 0.001$) and no effect of Definiteness nor an interaction (p's > 0.4). Therefore, our assumption is that the mention of the related noun, which contributes to the high context~response similarity in the high-related condition, likely reflects bridging, rather than coincidental use of a noun from the context sentence.

The proportions in **Table 6** also address another open question about our results. High context~response similarity can, in principle, be achieved without actually mentioning a contextually given high-related noun. For example, returning to our restaurant/waiter prompt,

both of the (constructed, not attested) responses shown in (19) have high similarity to the context sentence, yet only the (19a) response utilizes the high-related noun from the context sentence.

(19) Speaker A: Nigel and I went out last night to that new restaurant. A waiter was friendly and helpful.
Speaker B: Wait, sorry, I wasn't listening. Which waiter are you talking about?
a. Speaker A: A waiter at that new restaurant I went to with Nigel
b. Speaker A: A waiter that Nigel and I met when we went out last night

If responses to definites and indefinites systematically varied in this way, we would have clear evidence that definiteness is affecting the way in which participants bridge. However, **Table 6** shows that definiteness does not affect the rate of use of the high-related noun. So, whether the target NP is definite or indefinite, a high-related noun in the context sentence is equally likely to be used in constructing the response.

The analyses we present here represent a subset of our analyses; we also conducted several other exploratory analyses (e.g., computing scores for a bag of words to which no stemming/lemmatizing was applied or in which we eliminated the target noun). Across these analyses, the effect of Relatedness was robust, whereas the effect of Definiteness was not. We chose to report the current set of analyses, because we believe they represent the best treatment of the data for identifying similarity (stemming, lemmatizing) and because we wanted to be conservative and allow for the possibility of an independent effect of definiteness, if it is present, to be recognized.

In summary, then, this experiment suggests that entity relatedness is a significant predictor of a bridged interpretation, independent of definiteness (as shown in the context~response similarity scores); in particular, in the low-related condition, we do not see an increased rate of bridging of definites over indefinites, as predicted by the two-phenomena view. We, therefore, take our results to support the one-phenomenon view of bridging.

5. Experiment 2: Self-paced reading

This experiment uses a self-paced reading (SPR) task to further test for an independent effect of definiteness in bridging that would support the two-phenomena view. We assess reading times at the point in a sentence where an invited bridge is cancelled. To illustrate, passage (20) initially invites a bridged interpretation of *the window* as denoting a window in the living room mentioned in the context sentence. However, that interpretation is cancelled at the descriptive relative clause (RC) *that was in her dream*, which specifies that the window mentioned is *not* in fact the one in the living room.

(20) Jane was in the living room. The window that was in her dream suddenly came to mind.

We hypothesize that this cancellation of an invited bridge will lead to processing difficulty, which would be reflected in greater reading times. The more evidence there is – from the form of the NP or from the context – that a bridged interpretation is likely, the more the bridge-breaking RC content is expected to cause reading difficulty.

Passage (21) uses the same context sentence as (20), but the following sentence begins with the low-related noun *knife*. In this case, we hypothesize that the reader will be less likely to assign a bridged interpretation when reading *the knife*, and hence is less likely to experience processing difficulty when encountering the same relative clause, *that was in her dream*.

(21) Jane was in the living room. The knife that was in her dream suddenly came to mind.

As in Experiment 1, we manipulate both entity relatedness and definiteness. Again, of particular interest is the comparison between definite and indefinite NPs in the low-related condition, because only the two-phenomena view predicts a difference in behavior in that condition. As we will show, processing difficulty at the bridge-breaking RC is only evident in the high-related definite condition.

5.1 Participants

In order to create a dataset of 100 monolingual English-speaking participants, we recruited and paid 126 participants via Prolific ($2.50 for a task that was estimated to take 10 minutes). All participants indicated they were monolingual English-speaking US nationals, but 24 subsequently mentioned growing up with a non-English language at home when we asked about their language background. Those 24 were removed, as were a further 2 participants with low accuracy on the comprehension questions. With these exclusions, our dataset consists of 100 monolingual English participants' reading times.

5.2 Materials

The target items consisted of 40 passages that followed the structure of (20–21). As in Experiment 1, items included both singular and plural target nouns (28 singular, 12 plural). We varied the Relatedness and Definiteness of the target noun, as in the sample item sets in **Table 7**. Each item consists of a context sentence followed by a continuation sentence. The continuation sentence begins with a determiner-noun sequence followed by a restrictive relative clause. The determiner is either *the* or *a*. The following noun either denotes an entity that is highly related to an anchor entity introduced by the context sentence (*window/living room*) or one that is unrelated to the context sentence (*knife/living room*). As shown in the sample items, in the critical items, the determiner-noun sequence was always presented without the RC, to allow for an initial reading of this sequence as a complete NP.

Table 7 shows the chunking we used in the self-paced reading paradigm. The context sentence was presented as a single chunk and the continuation sentence had chunks for the determiner-noun sequence, the start of the RC, the bridge-relevant content of the RC, and one or more spillover regions.

Table 7: Experiment 2 materials for singular and plural NP nouns.

Singular NP target	
Context sentence: [Jane was in the living room.]$_{context}$	
[high, def]	[The window]$_{Det\text{-}N}$ [that was in]$_{RC}$ **[her dream]** $_{RC\_break}$ [suddenly came to mind.]$_{Spill}$
[high, indef]	[A window]$_{Det\text{-}N}$ [that was in]$_{RC}$ **[her dream]** $_{RC\_break}$ [suddenly came to mind.]$_{Spill}$
[low, def]	[The knife]$_{Det\text{-}N}$ [that was in]$_{RC}$ **[her dream]** $_{RC\_break}$ [suddenly came to mind.]$_{Spill}$
[low, indef]	[A knife]$_{Det\text{-}N}$ [that was in]$_{RC}$ **[her dream]** $_{RC\_break}$ [suddenly came to mind.]$_{Spill}$
Plural NP target	
Context sentence: [Barbara was grinning from ear to ear when she walked into her classroom.]$_{Context}$	
[high, def]	[The desks]$_{Det\_N}$ [that were in]$_{RC}$ **[a magazine]** $_{RC\_break}$ [that she had just seen]$_{Spill}$ [were exactly]$_{a}$ [what she wanted.]$_{b}$
[high, indef]	[Some desks]$_{Det\_N}$ [that were in]$_{RC}$ **[a magazine]** $_{RC\_break}$ [that she had just seen]$_{Spill}$ [were exactly]$_{a}$ [what she wanted.]$_{b}$
[low, def]	[The shoes]Det_N [that were in]$_{RC}$ **[a magazine]** $_{RC\_break}$ [that she had just seen]$_{Spill}$ [were exactly]$_{a}$ [what she wanted.]$_{b}$
[low, indef]	[Some shoes]$_{Det\_N}$ [that were in]$_{RC}$ **[a magazine]** $_{RC\_break}$ [that she had just seen]$_{Spill}$ [were exactly]$_{a}$ [what she wanted.]$_{b}$

The target region (in bold) consists of the portion of the RC that identifies the referent as a brand new item, unrelated to the context sentence. If a participant had already inferred a bridge to the context sentence (e.g., assumed, in the first example, that the window being mentioned was a window in the living room), then this target region would require them to revise that interpretation. For a given item, the context sentence, target, and spillover regions were the same for all conditions, and the only words that varied were the sentence-initial determiner and noun in the continuation sentence.

An additional 24 two-sentence passages were used as fillers. Of these, 8 followed the structure of the target items (context sentence, continuation-initial Det-Noun, RC), but they were set up to ensure that participants couldn't learn over the course of the experiment that potential NP bridges are always broken. In these fillers, the bridge-compatible noun at the start of the continuation sentence was always followed by a bridge-compatible RC. We also varied the chunking, to help prevent participants from associating the chunking of the target items with a determiner-noun sequence and then a bridge-breaking RC. In half of the 8 fillers with RCs,

the context sentence was presented in smaller chunks (to avoid Det-Noun-RC sequences only appearing after a single-chunk context sentence); in the other half, the continuation contained a Det-Noun-RC single chunk (to avoid a pattern where a single-chunk context sentence would always be followed by the separate Det-Noun and bridge-breaking RC). A further 16 fillers contained no RC. For the chunking, half of the no-RC fillers contained a single-chunk context sentence (to ensure that participants could not learn that single-chunk context sentences and sentence-initial Det-Nouns were always followed by an RC); the other half provided additional variability in the chunking patterns, by presenting a multi-chunk context sentence and a single-chunk continuation. See the Supplementary Materials for the full set of items.

5.3 Procedure

From Prolific, participants were directed to a website hosted by IbexFarm (Drummond, 2013) for the moving-window self-paced reading experiment. Passages initially appeared on the screen as a series of horizontal lines, where the line length corresponded to the length of the regions. Participants revealed each subsequent region of the passage by pressing the space bar on their keyboard. Passages were presented non-cumulatively, so that each newly revealed region was the only visible region on the screen.

After a quarter of the items, participants saw verification statements which they responded to by clicking on “TRUE” or “FALSE” with their cursor. They received feedback for incorrect answers only.

5.4 Results

Figure 3 and **Table 8** show the means of the raw reading times starting at the sentence-initial Det-Noun in the continuation sentence and proceeding through the RC, the RC_break, and the Spillover. Of note are the slow reading times at the bridge-breaking RC region for passages with a high-related definite NP. Contra the prediction of the two-phenomena view, definites and indefinites show no apparent difference in the low-related condition.

Table 8: Reading times, means and standard error (ms), by condition and region.

	Det-N	***RC***	***RC_break***	***Spillover***
high definite	758.61 ± 26.92	696.04 ± 26.86	1176.60 ± 49.78	673.89 ± 25.03
high indefinite	771.36 ± 28.14	676.35 ± 28.14	1074.69 ± 46.24	659.32 ± 21.73
low definite	784.06 ± 28.02	724.63 ± 28.49	1066.53 ± 42.88	654.69 ± 23.98
low indefinite	797.50 ± 33.51	717.70 ± 30.36	1083.01 ± 46.51	638.60 ± 20.69

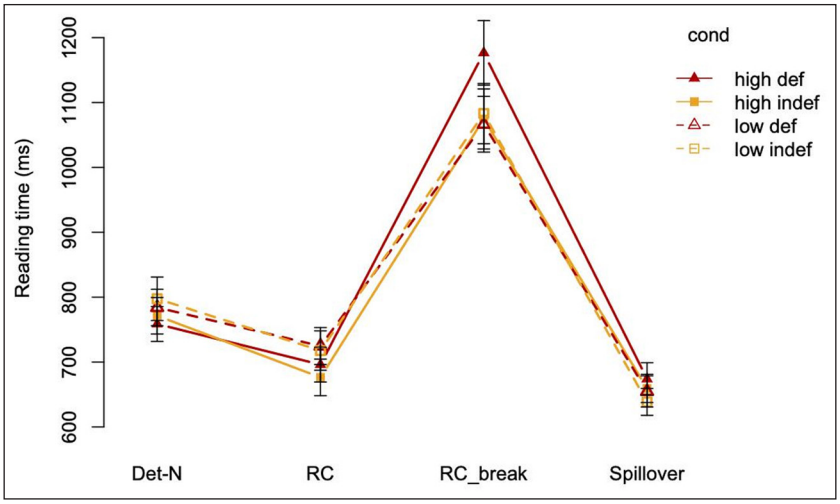

Figure 3: Experiment 2 reading times, calculated as by-participant means with standard error (ms) by condition, from Det-Noun region to the Spillover region.

For the analysis, we use linear mixed-effects regression models (LMER; Baayen et al., 2008), using the lme4 package in R (Bates et al., 2015; R Development Core Team, 2017). Traditional analysis of reading time data involves a series of separate analyses, one for each region. However, such an approach raises the possibility of a Type I error, given the non-independence of reading times at different positions in the same sentence and the use of multiple comparisons. Instead, we first build a single large model that contains fixed factors for Relatedness, Definiteness, and also Region. For interactions that reach significance in this large model, we conduct follow-up analyses. This latter approach limits the number of region-specific analyses by only targeting those whose interaction reached significance in the omnibus analysis. This follows recent work on multi-window analyses (see Grüter et al., 2018, for eye-tracking data, and Rohde et al., 2021, for self-paced reading). We use maximal random-effects structure as permitted by the data (see Barr et al., 2013).

Relatedness and Definiteness are coded as in Experiment 1 (–.5/+.5 for low/high and indefinite/definite), and Region uses the sentence-initial Det-Noun sequence as the reference level. Note that this coding means that an interaction between a particular region and one or more of the manipulated factors signals that the behavior of those factors at that region differs from their behavior at the Det-Noun. For example, the prediction that high relatedness supports bridging would correspond to an interaction between Relatedness and the RC_break region: At the Det-Noun region, high relatedness is expected to yield faster reading times (words that are semantically related to the preceding context are typically read faster), but if high relatedness supports a bridged interpretation, then at the RC_break region, high relatedness is expected to yield slower reading times when a posited bridge is broken.

Table 9 shows the model output for Experiment 2. The main effect of Relatedness indicates that high-related entities yield faster reading times at region's reference level, i.e., at the sentence-initial Det-Noun. The three main effects of region (Region_RC, Region_RC_break, Region_Spillover) indicate that the reading times in these regions differ significantly from that at the sentence-initial NP, independent of condition (faster at Region_RC, slower at Region_RC_break, and faster in Region_Spillover). As noted above, the role of Region is most relevant to our research question when Region interacts with one or more of the manipulated factors.

Table 9: Results of linear mixed-effects models of Experiment 2 reading time data. Boldface indicates significance.

	Beta	**SE**	**t**	**p**
(Intercept)	**801.43**	**38.44**	**20.85**	**<0.001**
Definiteness	−11.83	15.82	−0.75	0.45
Relatedness	**−31.61**	**15.82**	**−2.00**	**<0.05**
Region_RC	**−77.39**	**13.66**	**−5.67**	**<0.001**
Region_RC_break	**310.46**	**25.26**	**12.29**	**<0.001**
Region_Spillover	**−120.79**	**14.65**	**−8.24**	**<0.001**
Def:Rel	−21.36	31.66	−0.68	0.50
Def:Region_RC	24.71	22.37	1.11	0.27
Def:Region_RC_break	39.72	22.42	1.77	0.08
Def:Region_Spill	27.05	22.34	1.21	0.23
Rel:Region_RC	−8.57	22.37	−0.38	0.70
Rel:Region_RC_break	**78.90**	**22.41**	**3.52**	**<0.001**
Rel:Region_Spill	**56.26**	**22.33**	**2.52**	**<0.05**
Def:Rel:Region_RC	20.62	44.74	0.46	0.64
Def:Rel:Region_RC_break	**122.76**	**44.84**	**2.74**	**<0.01**
Def:Rel:Region_Spill	6.76	44.68	0.15	0.88

The 2-way Relatedness × Region_RC_break interaction is likely driven by the 3-way interaction at that region, which we describe in the paragraph below. In addition, we see a Relatedness × Region_Spillover interaction, whereby high relatedness yields reading times at the Spillover region that differ from those at the Det-Noun region. For this 2-way interaction, we conducted follow-up analyses. At the reference level Det-Noun region, we already know that high relatedness yields faster reading times than low relatedness (see **Table 9**: beta = –31.61,

SE = 15.82, t = –2.00, p < 0.05), whereas at the Spillover region, this difference is neutralized and the high-related condition instead yields marginally slower reading times than the low-related condition (beta = 20.33, SE = 11.60, t = 1.75, p = 0.09).

The 3-way Relatedness × Definiteness × Region_RC_break interaction can be understood to capture the slow mean reading time that is apparent in **Figure 3** in the high-related definite condition in the RC_break region. For this 3-way interaction, we conducted follow-up analyses. At the reference level Det-Noun region, we already know that the Relatedness × Definiteness interaction is not significant (**Table 9**: beta = –21.36, SE = 31.66, t = –0.68, p = 0.50), whereas at the RC_break region, the interaction is significant (beta = 85.12, SE = 41.11, t = 2.07, p < 0.05). A further follow-up of this interaction confirms that it is the high-related definite condition that yields the slowest reading time. Definites are slower than indefinites in the high-related condition (beta = 67.91, SE = 22.61, t = 3.004, p < 0.005), but there is no effect of Definiteness in the low-related condition (beta = –0.90, SE = 18.25, t = –0.05, p = 0.96).

As in Experiment 1, we conducted a series of post hoc analyses (see the Supplementary Materials) on the singular vs. plural subsets of the data and on the relational vs. non-relational subsets. The upshot is that some analyses show the same 3-way Relatedness × Definiteness × RC_Break interaction as in the main analysis, whereas some analyses (particularly in the smaller data subsets with potentially reduced power) show only a marginal 3-way interaction or simply a Relatedness × RC_Break interaction or a Relatedness × Spillover interaction. Crucially, we see no emergence of new effects or interactions with Definiteness.

The uniqueness rating study mentioned in relation to Experiment 1 (see Appendix C) included Experiment 2 items and showed that all but one of these items (*chair/leg*) met our criterion of plausibility of both unique and non-unique contextual instantiation. The chair/leg item, as noted, involves a robustly relational noun and, hence, is expected to be felicitous when used with a definite, despite contextual non-uniqueness.

5.5 Discussion

The strategy in this experiment is to use reading times at the RC_Break region as an indicator of the presence or absence of (early) bridging. Unlike Experiment 1, which provides evidence of an “all things considered” late-stage interpretation, Experiment 2 provides a window into incremental processing. The findings of this experiment reveal a more complex picture than Experiment 1. With respect to the low-related condition, Experiment 2 adds to the evidence of Experiment 1 that in a discourse situation providing no contextual support for bridging, definites are no more likely to bridge than indefinites. That condition provides evidence that this is not merely an “all things considered” effect, but shows that the initial encounter with a low-related definite does not generate an early attempt to bridge. Thus, there is no evidence that a definite drives a bridging attempt, independently of contextual support for such a reading. However,

in the high-related case, the situation is more complex. Although Experiment 1 indicates that interpreters are equally likely to give bridged interpretations to high-related definites and indefinites at a late stage in interpretation, things look different at the earlier processing stage targeted in Experiment 2. Only high-related definites show a slowdown in reading times at the RC_Break, indicating an early commitment to a bridged reading.

This interaction is unsurprising on the one-phenomenon view, which takes bridging to be a general interpretative response to multiple cues provided in the input. Arguably, it is also predicted on the two-phenomena view, because, as noted in Section 1, even those who take bridging to be a response to an apparently antecedentless definite also recognize that bridging is constrained by plausibility, so such definites with an "obvious" anchor for bridging will be more likely to bridge than definites with no such anchor. However, the idea of bridging as definiteness-driven suggests that we are more likely to see an early effect of definiteness, regardless of relatedness, with bridging attempts abandoned as broader coherence considerations come into play in later interpretation. Our data show no such effect, and this is consistent with results from more fine-grained experimental paradigms which we discuss in the following section.

6. General discussion: One phenomenon or two?

We began this article by considering a pair of examples (repeated here) which appear superficially entirely parallel. In both cases, the subject NP *the door/a door* is most naturally understood to refer to a door of the just-mentioned house. We call this a bridged reading of the NP:

(1) a. Yasmin approached the house. The door was open.
b. Yasmin approached the house. A door was open.

Most of the theorizing about bridged interpretations of NPs has focused on bridging of definites, as in (1a). These theories of bridging have understood the phenomenon as involving a response on the part of the hearer to the occurrence of an antecedentless definite. This posited response aims to resolve an apparent violation of the contextual requirements of the definite. The attempt at resolution may be foiled when the context does not provide an appropriate anchor for bridging, as in (2):

(2) Yasmin approached the house. The store was open.

And the attempt is facilitated when an "obvious" anchor is provided, as in (1a). But, on this view of bridging, it is definiteness that is the principal driver of bridging, rather than knowledge of typical relations between entities or other considerations of coherence. This understanding of the bridging of definites requires the bridging of indefinites to be explained in a different way, giving rise to the view that bridging of definites and bridging of indefinites constitute two different phenomena.

The alternative view is that the similarity between (1a) and (1b) is not merely superficial but substantive, that both are instances of a single phenomenon. This view does not deny that definites (and indefinites) impose constraints relevant to interpretation. A speaker's choice to introduce a new referent with a definite or with an indefinite carries information about the intended referent, such as whether or not it is contextually unique; this information can either support or conflict with coherence-driven expectations. But, on a one-phenomenon view, definiteness cannot be the principal driver of bridging.

Distinguishing the views empirically is challenging, because, as we have noted repeatedly, both views predict that interpreters are sensitive to both entity relatedness (coherence) and definiteness. We have identified the contexts *lacking* an obvious anchor for a newly introduced entity as the central testing ground of the two views. If bridging of definites is primarily driven by definiteness, but bridging of indefinites is driven solely by coherence considerations, then interpreters should be more inclined to invent bridged interpretations for definites than for indefinites in the absence of a bridge-inviting anchor.

We have, therefore, tested rates of bridging, varying two factors: linguistic marking of definiteness and the presence or absence of a high-related noun in the immediate context. Our experiments explore the rates of bridging with two different methodologies, allowing us to investigate bridging judgements at two different timescales. Both experiments are revealing, and point to interesting differences between different stages of interpretation.

Experiment 1, utilizing a dialogue continuation task, explores the "all things considered" interpretation of a target NP when both the context sentence and the continuation sentence containing the target are fully available. The experiment found that in contexts with a high-related anchor (high-related condition) and in contexts with no such anchor (low-related condition), rates of bridging of definite and indefinite target NPs were not significantly different. The only predictor of bridging in this experiment was the presence or absence of a high-related anchor.

Experiment 2, a self-paced reading task, investigates NP interpretation at an earlier point in sentence processing. The critical region lies within a relative clause modifier of the target noun, at a moment when it becomes apparent to the reader (participant) that the noun is *not* intended to be bridged. By measuring processing difficulty at this point (operationalized by reading time), we assess whether or not the participant had adopted a bridged interpretation in the course of sentence interpretation, very soon after encountering the target noun. The results of this experiment were more nuanced. We found a significant interaction between definiteness and the availability of a high-related anchor, such that only definites in the high-related condition showed evidence of a processing slowdown at the bridge-incompatible information. There were no significant differences among the remaining conditions.

On the one hand, the lack of difference between definites and indefinites in the low-related condition again militates against the two-phenomena view. There is no evidence from the reading times in the low-related condition that participants who have just read a definite NP without an obvious anchor have higher expectations of a bridged reading than participants who have just read an indefinite NP. On the other hand, a purely coherence-driven view would predict that both definites *and* indefinites in the high-related condition would give rise to the same expectation of bridging, and, hence, that the bridge-incompatible information would cause the same processing difficulty in both cases. This is not, however, what we found.

The evidence suggests that at the stage of processing we target, participants have not committed to a bridged interpretation *except* when the presence of a high-related anchor provides a natural solution to the problem of satisfying the requirements of a definite. This is consistent with a one-phenomenon view, recognizing that overall interpretation will not only reflect an expectation of coherence, but must also render the speaker's use of a definite (or indefinite) felicitous in the context. However, the results, overall, could also be seen to be compatible with a two-phenomena view incorporating an early role for *rejection* of bridging in the absence of an anchor. It seems to us, however, that the most plausible explanation of the results from the two experiments is that interpreters encountering a newly introduced referent bring to bear both entity relatedness (coherence considerations) and any signals provided by definiteness, with neither being the primary driver in either the definite or indefinite case.

Our SPR task provides an interesting intermediate timescale between the ERP experiments of Schumacher (2009), discussed in 2.2., and our Experiment 1. Schumacher's neurological measure provides evidence of moment-by-moment responses to linguistic input, looking at responses within the range of 400–600 milliseconds after encountering the target noun. Schumacher argues that in early interpretation, hearers attempt to integrate new NPs to existing representations regardless of definiteness, and that definiteness shows a measurable effect only with respect to the P600. Given that in our SPR task, it takes participants approximately 750–820 ms to read the Det-N sequences, our presentation of bridge-incompatible information comes considerably later than the onset of the definiteness effect found by Schumacher. Schumacher's data suggest that indeed, by this point, both the relatedness of the new entity to the context and information from definiteness marking have been incorporated into interpretation.

However, these early measures do not show us how interpreters ultimately integrate the multiple sources of information provided by context, content and sentence form to arrive at a final interpretation of an utterance. This is precisely the information provided by Experiment 1. These results suggest that ultimately, the relatedness of a new entity to prior context is an overriding factor in the decision to bridge or not. These observations raise interesting questions about the time course of integration of linguistic signals and general considerations of coherence. As noted, Schumacher suggests that the search for coherence applies early, with the morphosyntactic effects

of definiteness coming in later. Our results, to put it very informally, suggest that plausibility is ultimately the decider with respect to the final preferred interpretation. One question here is whether the integration effects that Schumacher observes are of a kind with the perhaps more explicit considerations of plausibility that we typically think about when assessing coherence.

With respect to the one-phenomenon-or-two question that we have focused on here, the answer is likely to turn out to be complex. It may turn out that there is a stage of processing – the stage targeted by Experiment 2 – at which definites and indefinites are, indeed, being treated differently, but that at both very early stages and late stages, the processing system (perhaps for different reasons) is somewhat indifferent to linguistic cues. We hope that further work will continue to explore this question. Both of the experiments reported here contain new methodology which, to our knowledge, is the first which operationalizes bridging and provides specific measures to assess whether bridging has occurred, either in the course of interpretation (Experiment 2) or at its conclusion (Experiment 1). Prior experimental work on bridging has tended to assume that in straightforward examples like (1a-b), bridging occurs. However, in order to experimentally investigate factors which contribute to bridging, it is essential to have objective measures of the presence of bridging. We offer our methodologies as a contribution to the empirical study of bridging.

Appendix A: Experiment 1 elicitation task

We used the participant responses in the elicitation task to construct materials for the main task in Experiment 1. Here, we give more background on the setup and motivation for the task, which was initially intended to measure bridging itself, but whose challenges led to the development of the "Which N?" task in Experiment 1. We also provide the instructions and full list of target prompts in the Supplementary Materials.

In the elicitation task, participants were presented with a context sentence followed by an NP (either definite or indefinite) and were asked to complete the sentence. This task was originally designed to assess rates of bridging of the presented NP in the absence of further cues to passage coherence, with the expectation that the continuation provided would make clear whether a bridged or nonbridged interpretation had been adopted. In some cases, it was easy to determine whether bridging was present or absent. In (i), *some buttons* must almost certainly refer to buttons belonging to the shirt, given the continuation *falling off of it*. In (ii), *a chair* can only plausibly refer to some chair unrelated to the fruit arrangement or to Hilda's creation of it, given no visible link to the context sentence.

(i) I'm really annoyed about this new shirt. Some buttons … are already falling off of it.
(ii) Hilda created a nice arrangement of fruit. A chair … was on the porch.

However, in most cases, the continuation left it unclear whether or not the NP prompt was being linked to an entity in the context sentence. In some cases, the same continuation appeared in multiple conditions (iii–iv).

(iii) a. Jane was in the living room. A window was open.
b. My best friend had the most beautiful wedding. The window was open.

(iv) a. Jane was in the living room. A window is very good.
b. Tania took the kids to the playground. A slide is very good.
c. Barbara was grinning from ear to ear when she walked into her classroom. The fish is very good.

It became apparent that establishing consistent criteria for coding bridging would be difficult and that our annotation of bridging merely reflected the interpretation that the annotator found most natural for the completed string. To evaluate the most likely interpretation of the NPs in the continuations provided, we therefore developed the study that we report in Experiment 1, where we asked a separate set of participants to read a sampling of the elicited continuations and indicate for us how they would interpret the NP.

Appendix B: Experiment 1 similarity scores for sample participant responses

Listed below are a sample of participant responses (underlined) from Experiment 1 that illustrate the range of responses and their corresponding context~response and continuation~response similarity scores. Stemmed similarity and lemmatized similarity scores are separated by /. Speaker B's apology and question are abbreviated to show only the question with the target noun.

(i) A: Nigel and I went out last night to that new restaurant. A waiter was very friendly.
B: …Which waiter?
A: The one at the new restaurant Nigel went to last night
[context~response: 0.92/0.92; continuation~response: 0.00/0.00]

(ii) A: Nigel and I went out last night to that new restaurant. A waiter was friendly and helpful.
B: …Which waiter?
A: A waiter at that new restaurant I went to last night with Nigel.
[context~response: 0.93/0.93; continuation~response: 0.22/0.22]

(iii) A: Paula is going to the zoo. Some professors are going as well.
B: …Which professors?
A: Some professors that are going to the zoo with Paula
[context~response: 0.87/0.87; continuation~response: 0.58/0.58]

(iv) A: Michael and Ray were playing their new board game. A lens was required to play it properly.
B: …Which lens?
A: The lens is to help Michael and Ray play their new board game properly.
[context~response: 0.82/0.68; continuation~response: 0.50/0.50]

(v) A: I want to try to fix this old chair. The rug is under it.
B: …Which rug?
A: The rug under the old chair.
[context~response: 0.58/0.58; continuation~response: 0.52/0.52]

(vi) A: Charles has a large table top aquarium. The fish fit perfectly in it.
B: …Which fish?
A: Charles's fish!
[context~response: 0.00/0.32; continuation~response: 0.41/0.41]

(vii) A: Rosie was putting the groceries away in the cupboard. The shelf was full of goodies for later.
B: …Which shelf?
A: The shelves that are used to store the groceries
[context~response: 0.22/0.22; continuation~response: 0.00/0.25]

(viii) A: Yesterday, Tony finally went to see his dentist. The hygienist raved about his dental care.
B: …Which hygienist?
A: The hygienist at the dental clinic Tony visited.
[context~response: 0.18/0.18; continuation~response: 0.45/0.45]

(ix) A: Hilda created a nice arrangement of fruit. The chair was not made of fruit.
B: …Which chair?
A: A chair in the living room. It wasn't made of fruit.
[context~response: 0.18/0.18; continuation~response: 0.71/0.71]

(x) A: I want to do something about my living room floor. A leg damaged the old flooring.
B: …Which leg?
A: The leg that damaged the old flooring
[context~response: 0.22/0.00; continuation~response: 1.00/1.00]

Appendix C: Uniqueness norming study

Our norming study was similar to that reported in Singh et al. (2016), and assessed the likelihood of unique versus multiple instantiation of the target nouns in our item contexts. We asked participants to judge the likelihood of both unique instantiation in the context (*Imagine someone <item context>, how likely is it that there is exactly one <noun>?*) and multiple instantiation (*Imagine someone <item context>, how likely is it that there are multiple <noun>?*), manipulated within items and within participants. The question phrasing varied across items to match the verb and anchor noun in each item's context sentence (e.g., *Imagine someone going to a stable, Imagine someone has glasses, Imagine someone coasting on a bicycle, Imagine someone staying late in an office*). We followed Singh et al. (2016) in adjusting the phrasing of a subset of items to highlight a primary encounter within a larger setting (e.g., *Imagine someone has just been admitted to hospital, how likely is it that they encounter exactly one doctor/multiple doctors?*).

We tested all Experiment 1 and Experiment 2 items that contained a singular target noun (33 total, 21 common across both experiments, 5 used only in Experiment 1 and 7 used only in Experiment 2). Alongside the target items, we included 8 control items describing entities that either must be contextually unique to be plausible (e.g., *Imagine there's a one-man show, how likely is it that there's exactly one performer?*) or must be contextually non-unique (e.g. *Imagine there's a bird flying overhead, how likely is it that it has exactly one wing?*), counterbalanced to elicit likely/unlikely ratings for questions about exactly one or multiple entities. Item order was randomized for each participant, so it was unlikely that a participant saw the unique and multiple conditions of a given item one after another. The task used a slider bar (0–100) with 'unlikely' and 'likely' at the endpoints. Twenty native speakers of English in the UK were recruited from Prolific and paid £4 for a task that took less than ten minutes. All participants performed well

on the control items, with plausible controls receiving a mean rating of 94.43 and implausible controls receiving a mean rating of 4.44. For the target items, the mean norming ratings for the exactly one and multiple questions were 42.41 and 61.65, respectively, both much higher than the implausible controls, with the exception of two items whose ratings in the 'unique' condition failed to distinguish them from the unique-implausible fillers. For completeness, we re-ran the analyses, excluding these two items. As expected by excluding only two items, the results were unchanged. See the Supplementary Materials for the full details of the norming study, including the by-item ratings.

Data accessibility statement

All datasets and scripts for statistical analysis can be accessed here: https://osf.io/6w7v8/?view_only=dbbccc0a5d83482eb998bd17e539b409.

Ethics and consent

All studies were reviewed and approved by the ethics boards of the University of Edinburgh's School of Philosophy, Psychology, and Language Sciences and by the Carnegie Mellon University Internal Review Board (Study # STUDY2020_00000019).

Acknowledgments

We thank Dustin Updyke for work on the computation of similarity scores for Experiment 1. This work was supported in part by a Leverhulme Trust Prize in Languages and Literatures to H. Rohde, and research funding provided to M. Simons by the CMU Department of Philosophy.

Competing interests

The authors have no competing interests to declare.

Author contributions

M. Simons and H. Rohde contributed to: conceptualization, funding acquisition, investigation, methodology, writing original draft, review, and editing. H. Rohde was responsible for formal analysis and visualization.

References

Asher, N., & Lascarides, A. (1998). Bridging. *Journal of Semantics, 15*, 83–113. https://doi.org/10.1093/jos/15.1.83

Asher, N., & Lascarides, A. (2003). *Logics of conversation*. Cambridge University Press.

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Barker, C. (2005). Possessive weak definites. In J.-Y. Kim, Y. Lander, & B. H. Partee (Eds.), *Possessives and beyond: Semantics and syntax* (pp. 89–113). GLSA Publications.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, *68*, 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*, 1–67. https://doi.org/10.18637/jss.v067.i01

Bird, S., Loper, E., & Klein, E. (2009). *Natural language processing with Python*. O'Reilly Media Inc.

Clark, H. H. (1973). *Comprehension and the Given-New contract* [Conference presentation]. Conference on the Role of Grammar in Interdisciplinary Linguistic Research, University of Bielefeld, Bielefeld, Germany.

Clark, H. H. (1975). Bridging. In R. C. Schank & B. L. Nash-Webber (Eds.), *Theoretical issues in natural language processing.* (pp. 175–196). Association for Computing Machinery. Reprinted in P. N. Johnson-Laird, & P. C. Wason (Eds.), *Thinking: Readings in cognitive science* (pp. 411–420). Cambridge University Press. 1977.

Clark, H. H., & Haviland, S. E. (1977). Comprehension and the given new-contract. In R. O. Freedle (Ed.), *Discourse processes: Advances in research and theory: Discourse production and comprehension* (pp.1–40). Ablex Publishing.

Clifton, C. (2012). Situational context affects definiteness preferences: Accommodation of presuppositions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(2), 487–501. https://doi.org/10.1037/a0028975

Coppock, E. (2022). On definite descriptions: Can familiarity and uniqueness be distinguished? In D. Altshuler (Ed.), *Linguistics meets philosophy* (pp. 109–136). Cambridge University Press. https://doi.org/10.1017/9781108766401.007

Drummond, A. (2013). Ibex Farm. http://spellout.net/ibexfarm/.

Frazier, L. (2006). *The big fish in a small pond: Accommodation and the processing of novel definites* [Manuscript]. University of Massachusetts, Amherst.

Grüter, T., Takeda, A., Rohde, H., & Schafer, A. (2018). Intersentential coreference expectations reflect mental models of events. *Cognition*, *177*, 172–176. https://doi.org/10.1016/j.cognition.2018.04.015

Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior, 13*, 512–521. https://doi.org/10.1016/S0022-5371(74)80003-4

Hawkins, J. A. (1978). *Definiteness and indefiniteness: A study in reference and grammaticality prediction.* Humanities Press.

Heim, I. (1982). *The semantics of definite and indefinite noun phrases* [Doctoral dissertation]. University of Massachusetts at Amherst. https://semanticsarchive.net/Archive/jA2YTJmN/Heim%20Dissertation%20with%20Hyperlinks.pdf

Hobbs, J. R. (1979). Coherence and coreference. *Cognitive Science*, *3*, 67–90. https://doi.org/10.1207/s15516709cog0301_4

Kehler, A. (2015). Reference in discourse. In S. Lappin & C. Fox (Eds.), *The handbook of contemporary semantic theory* (2nd edition). John Wiley and Sons. https://doi.org/10.1002/9781118882139.ch20

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, *8*(1), 339–359. https://doi.org/10.1007/BF00258436

Li, J., & Ettinger, A. (2022). Heuristic interpretation as rational inference: A computational model of the N400 and P600 in language processing. *Cognition*, *233*, 105359. https://doi.org/10.1016/j.cognition.2022.105359

Link, G. (1983). The logical analysis of plurals and mass terms: A lattice-theoretical approach. In R. Bäuerle, C. Schwarze, & A. V. Stechow (Eds.), *Meaning, use, and interpretation of language* (pp. 302–323). De Gruyter. https://doi.org/10.1515/9783110852820.302

Manning, C., & Schütze, H. (1999). *Foundations of statistical natural language processing*. MIT Press.

Prince, E. (1992). The ZPG letter: Subjects, definiteness, and information-status. In W. C. Mann and S. A. Thompson (Eds.), *Discourse description: Diverse linguistic analyses of a fund-raising text*. (pp. 295–325). John Benjamins. https://doi.org/10.1075/pbns.16.12pri

R Development Core Team. (2017). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. ISBN 3-900051-07-0.

Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, *26*(3), 287–350. https://doi.org/10.1023/A:1024157132393

Rohde, H., Futrell, R., & Lucas, C. G. (2021). What's new? A comprehension bias in favor of informativity. *Cognition*, *209*, 104491. https://doi.org/10.1016/j.cognition.2020.104491

Schumacher, P. B. (2009). Definiteness marking shows late effects during discourse processing: Evidence from ERPs. In S. L. Devi, A. Branco, & R. Mitkov (Eds.), *Anaphora processing and applications: Lecture notes in artificial intelligence* (Vol. 5847, pp. 91–106). Springer. https://doi.org/10.1007/978-3-642-04975-0_8

Schwarz, F. (2009). *Two types of definites in natural language*. [Doctoral dissertation, University of Massachusetts at Amherst]. GLSA Publications. https://doi.org/10.7275/1077454

Sharvy, R. (1980). A more general theory of definite descriptions. *The Philosophical Review*, *89*(4), 607–624. https://doi.org/10.2307/2184738

Singh, R., Fedorenko, E., Mahowald, K., & Gibson, E. (2016). Accommodating presuppositions is inappropriate in implausible contexts. *Cognitive Science*, *40*(3), 607–634. https://doi.org/10.1111/cogs.12260

van der Sandt, R. A. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, *9*(4), 333–377. https://doi.org/10.1093/jos/9.4.333