False but slow:
Evaluating statements with non-referring definites

Florian Schwarz
University of Pennsylvania

October 30, 2014

Abstract

One central debate in the analysis of definite descriptions concerns the truth-value of sentences where there is no entity that meets the description in the definite. Classical Russellian accounts predict them to be plain false, whereas presuppositional accounts predict them to be infelicitous. Recent discussions have homed in on the factors that affect actual judgment behavior in relation to the underlying status posited by different accounts. This paper presents experimental evidence for a presuppositional view based on response times for judging statements with non-referring definites to be 'false', which were longer relative to control statements where existence was asserted. I discuss the theoretical implications of these results, as well as of other findings from the literature, arguing that they support a presuppositional view of definites that sees the existence presupposition as conventionally encoded. The paper also makes a methodological contribution, as systematic evidence on speakers' judgments in these cases turns out to be hard to come by. Finally, the results inform the more general issue of the online processes involved in the interpretation of presupposed, as opposed to asserted, content.

1 Introduction

One central debate in the analysis of definite descriptions concerns the truth-value of sentences where there is no entity that meets the description in the definite. Classical Russellian accounts predict them to be plain 'false', whereas presuppositional accounts predict them to suffer from presupposition failure, which leads to infelicity. While most people seem to share a sense of 'squeamishness' (Strawson, 1964) about sentences such as *The king of France is bald*, it turns out to be surprisingly hard to get systematic evidence on speakers' judgments even on such simple cases, and little has been done to gather such evidence. In one recent attempt, Abrusán & Szendrői (2013) presented speakers with a variety of sentences containing non-referring definites and gave them an option of responding with 'can't say', in addition to 'true'

and 'false' (truth of the statements was based on world knowledge, which was independently controlled for). Between 80 and 90% of the time, subjects responded 'false' to statements that should have been truth-valueless, and thus triggered a 'can't say' response, on a presuppositional account, and the rate of 'false' responses for these statements did not differ from control sentences that were plain 'false' based on their truth-conditional content.[1] Given that linguists' and philosophers' judgments about presupposition failure, together with observations about presupposition projection, form the core of the notion of presupposition more generally, these difficulties in finding systematic evidence on the matter seem worrisome.

To address such worries, the experiments reported here take a different approach to gathering systematic evidence on rejecting statements with non-referring definites, by looking at the time-course of responses to shed light on the processes involved in reaching a 'false'-judgment. We look at cases where speakers seem to have relatively few qualms about judging a sentence with a non-referring definite as 'false', and the task employed does not even provide a third option. This provides a different angle on the question of whether there are differences between sentences that are plain false due to their literal meaning and ones that suffer from the existence requirement of a definite contained in them not being met: rather than striving to find potential differences in judgments, we strive to find differences in how judgments that are one and the same on the surface are arrived at. The contrast in predictions between a classical Russellian account and a standard presuppositional account remains the same: the former sees 'false'-judgments in both cases as entirely on par in that they are based on truth-conditional content. The latter, on the other hand, would assume that there are differences between the two cases, so that the paths to the 'false' judgments should differ. This would of course be supported by finding differences in behavioral response variables, and the experiments reported here yield evidence precisely along these lines.

One point that presuppositional theories have to account for is the fact that a sentence (or utterance) suffering from presupposition failure can lead to a 'false'-judgment by a speaker in the first place. This was already acknowledged by Strawson, who discussed speakers' varying inclination to judge a sentence with a non-referring definite as 'false', as opposed to expressing the inability to make a judgment, based on various properties of the examples at hand. For example, speakers seem to judge *The king of France visited the Exhibition* as false more readily than *The king of France is bald*. The mapping between naive speakers' truth-value judgments about the relevant sentences and the status assigned to them in theoretical terms can't be an entirely straightforward one, then. In fact, there is an ongoing debate concerning the question of precisely what factors are at play in this regard; see, for example Lasersohn (1993); von Fintel (2004); Schoubye (2010); Abrusán & Szendrői (2013), which are briefly reviewed below. Given this literature, there

[1] But note that they found interesting differences in judgments for negated versions of the sentences; see discussion below.

don't seem to be any principled obstacles to accommodating the existence of such variation in judgments within a presuppositional perspective (for discussion, see, for example, von Fintel, 2004; Schoubye, 2010). One option is that it seems perfectly reasonable to say that speakers can interpret the 'false' choice in the relevant types of tasks in terms of a more general rejection of the utterance, where the grounds for rejection could include pragmatic inappropriateness, e.g., based on presupposition failure. Alternatively, certain cases of non-referring definites may provide grounds for reaching genuine 'false'-judgments in a way different from 'false'-judgments based on truth-conditional content. However the presuppositional theorist chooses to spell this out, it will generally hold that something different may be going on in reaching a 'false'-judgment when this is due to presupposition failure as opposed to the falsity of truth-conditional content. In contrast, a simple Russellian account has no grounds for positing such differences. Any differences in the time course of reaching a 'false'-judgment in the two cases would therefore seem to provide support for a presuppositional view, as there would be a straightforward correlation between a posited theoretical difference and a measured empirical difference. The experimental evidence presented here will reveal such an empirical difference.

While the data speak against a classical Russellian account, we need to consider variations thereof that incorporate presuppositionality in one way or another. Leaving details to the general discussion section, I argue that the present results speak against accounts based on pragmatic presuppositions, and furthermore spell out the properties that a version of a neo-Russellian account would have to have in order to account for the results. The argument in favor of a conventionally encoded presupposition has important implications for the renewed debate in presupposition theory about whether any presuppositions are encoded conventionally. In connection with this, it will also be interesting to relate the present data to other recent work on local accommodation of presuppositions (Chemla & Bott, 2013), as it provides a potential alternative perspective on the data that may be particularly helpful for such neo-Russellian accounts.

While we may not be able to decisively differentiate all analyses of definite descriptions based on the present experiments, there nonetheless is an important methodological contribution that opens the path forward to further experimental work: given the difficulties of getting systematic judgment evidence, it turns out to be fruitful to investigate other response variables, such as response times, to contribute to the empirical body of evidence on which we base our theory.[2] Finally, to the extent that we accept that the present results support a presuppositional view of definites, we can also consider them as informing our more general understanding of how presuppositions are interpreted in online processing, which is a topic whose experimental investigation has only begun very recently, and where much remains to be learned.

[2]Similar approaches have been highly productive in other areas of Experimental Pragmatics, e.g., in the study of implicatures following Bott & Noveck (2004).

The paper is structured as follows. After briefly introducing the two basic theoretical perspectives on the semantics of definites under consideration, I review some of the existing empirical evidence and proposed accounts for variation in truth-value judgment behavior. I also highlight some of the relevant work in experimental psychology that provides a background for the present experiment and which, for the most part, has not been taken into consideration in theoretical debates in linguistics and philosophy. I then present the experimental design and the results from the present studies and discuss their implications for theories of definites and their relation to empirical truth-value judgment data.

2 The meaning of definites and truth value judgments

2.1 Definite descriptions and truth-values

Russell (1905) presents a quantificational view of the definite article, according to which it introduces both an existence and a uniqueness requirement.[3] The meaning of the sentence in (1a) then can be rendered as in (1b):

(1) a. The king of France is bald.
b. $\exists x\ [KoF(x)\ \&\ \forall y[KoF(y) \rightarrow x = y]\ \&\ B(x)]$

Put into words, an assertion of (1a) ends up making three separate claims: i) that there is a king of France, ii) that there is no more than one king of France, and iii) that the king of France is bald. Since the individual open statements are combined by conjunction, the sentence as a whole will be false if any of the corresponding properties fails to be instantiated. And since there is no present king of France, the sentence in (1a) is predicted to be false.

In line with seminal remarks by Frege (1892), Strawson (1950, 1964) argued against Russell's proposal by claiming that definite descriptions are referential, and furthermore give rise to a presupposition that there is an entity that uniquely satisfies the description introduced by the noun phrase. For Strawson, this rendered utterances of sentences like (1a) without a truth-value. Many different versions of presuppositional accounts of definite descriptions have been spelled out subsequently. Some locate presuppositions on the level of the semantics, so that sentences (rather than utterances thereof) end up without a truth-value, while others (e.g., Stalnaker, 1974) maintain a primarily pragmatic view, according to which it is speakers that presuppose (this position is compatible with, but does not require, encoding presuppositions semantically). There also are a number of technical options for handling the lack of a truth-value of a sentence, e.g., by defining definites in terms of partial functions, or by introducing a third truth-value, as on trivalent theories. Until picking up the issue of (purely) pragmatic presuppositions in the

[3] My discussion will focus almost exclusively on the existence requirement, since it is the one that is tested in the experiments described below.

general discussion, my discussion of a presuppositional view will assume that presuppositions are part of what is encoded conventionally as part of the lexical entry of expressions such as *the*, and that definite descriptions are referential, i.e., semantically of type e. The rendering of such an analysis in the format of Heim & Kratzer (1998) is representative of such a view (also see Schwarz, 2009; Elbourne, 2013, for detailed discussion of this approach within a situation semantic framework).[4]

(2) $[\![\textbf{the}]\!] = \lambda f : f \in D_{\langle e,t\rangle}$ and there is exactly one x such that $f(x) = 1$. the unique y such that $f(y) = 1$.

One of the central facts about presuppositions is that they, unlike asserted content, remain part of what is conveyed by the entire sentence even when the expression giving rise to them occurs with sentential negation. Accordingly,

(3) The king of France is not bald.

still presupposes that there is a unique king of France, even though the negation here is commonly seen as applying to the entire sentence *The king of France is bald*.[5] While this is a hallmark phenomenon from the perspective of presuppositional accounts, note that Russellian accounts can also deal with this fact, as they can assume such examples to involve a scopal ambiguity between the definite and negation. When the definite takes wide scope, the existence requirement is present at the level of the entire sentence.[6] While much has been written on these issues, this shall suffice as a brief survey of the lay of the land for present purposes.

Turning to the issue of how truth-value judgments are affected when the existence requirement is not met, Strawson's initial observation was that speakers feel 'squeamish' about sentences such as (1a). However, he himself also observed already that some non-referring definites seem to give rise to straightforward 'false'-judgments, e.g., in (4):

(4) The exhibition was visited yesterday by the king of France.

(Strawson, 1964)

[4]Following Heim and Kratzer, I place the presupposed content between the colon and the period; the asserted content follows the period.

[5]Parallel examples with the definite in the scope of negation in the surface form can easily be constructed to make the same point, e.g., *John didn't meet the king of France.*

[6]In contrast, the alternative narrow scope interpretation predicts the existence requirement to be absent. And in fact, it seems possible to find cases where the existence requirement does end up negated, e.g., with continuations such as the following:

1. The king of France is not bald - because there is no king of France!

While this is entirely expected from a Russellian perspective, presuppositional accounts have to add something to their story to capture these latter cases, e.g., by claiming that presuppositions can be interpreted 'locally' in certain circumstances (Heim, 1983), or that there are variants of negation that differ in their relation to presuppositional expressions (e.g., meta-linguistic negation Horn, 1985).

He suggested that the presence of the presupposition might be related to whether or not the definite serves as the topic of the sentence (see also Reinhart, 1981). Later accounts, such as Lasersohn (1993) and von Fintel (2004) focused on notions of verifiability, arguing that 'false'-judgments result just in case the hearer is able to verify the falsity of the claim in question if he temporarily assumes that the existence presupposition is in fact met. Most recently, Schoubye (2010) has laid out a novel proposal that makes the issue of squeamishness vs. falsity turn on the pertinence of the relevant statement to the Question under Discussion (QUD; Roberts, 1996), which accounts for a great deal of contextual variation in judgments. While the issue of which of these proposals ultimately is most successful at capturing the facts is an important one, the main line of argument in the present paper is mostly orthogonal to the differences between them. What it aims to establish is simply that the process of reaching a 'false' judgment differs between truth-conditionally false cases and ones with non-referring definites. To the extent that the proposals for predicting when precisely squeamishness arises are compatible with that, we will not attempt to further adjudicate between them. As a final remark, let us note that Russellians face a similar challenge, but also have options for meeting it: their default prediction is that the relevant sentences should simply be false, and they have to come up with an explanation of the cases where we feel squeamish. See Schoubye (2010) for some discussion in this direction, though he ultimately takes his proposal to provide at least tentative evidence against the options that he sees for Russellians.

2.2 Experimental work on the effects of non-referring definites

While the psycho-linguistic literature contains a host of studies on the processing of definite descriptions, little work has focused on cases where the existence requirement is clearly false. Many studies take an anaphoric or familiarity-based view of definites and manipulate the presence or absence of antecedents, the availability of bridging interpretations, or potential processing costs of accommodation (in the sense of Lewis (1979); see Murphy, 1984; Garrod & Sanford, 1982; Burkhardt, 2006, among many others).[7] While the last case comes close to the issue we're interested in, we want to study the effects resulting from the existence requirement being clearly false in a given context, rather than of mere lack of contextual support. There is one recent experimental study within linguistics that tries to get at this issue, namely Abrusán & Szendrői (2013).[8] In addition, there are a number of older studies from the psychology of memory, which deserve attention in the present debate, as they provide direct empirical evidence for the presuppositionality of the

[7] Another line of work has targeted the effect of the uniqueness requirement on online processing, e.g., Altmann & Steedman (1988); Altmann (1998); Crain & Steedman (1985).

[8] As an anonymous reviewer points out, recent work by Singh et al. (2013), which looks at effects of plausibility on accommodation, may also be relevant to the larger perspective under consideration here.

existence requirement of definites.

Abrusán & Szendrői (2013) aim to test several accounts of the variation in judgments between squeamishness and falsity by presenting sentences with definites and asking subjects to judge these as 'true' or 'false', with a third option of 'can't say'. As mentioned above, the positive versions of their sentences did not render any significant results. Subjects were generally quite happy to judge these as 'false' (at rates of 80-90%). However, the results were rather different for negated versions of the sentences. Note that in such cases, judgments in the absence of squeamishness should lead to 'true', rather than 'false', responses. Based on variations of both topicality and verifiability, Abrusan and Szendroi report differences in the proportions of true responses, ranging from 20% for the baseline condition ('The king of France is not bald') to 70% for a claim that involves an independently known individual ('The king of France is not married to Carla Bruni'), which is intended as a test of the proposal by (Lasersohn, 1993). While their results arguably lend support to both von Fintel's (1994) and Lasersohn's (1993) proposals in terms of the factors that affect truth-value judgments with non-referring definites, they do not necessarily provide specific support to any of the theoretical perspectives on definites. As Abrusan and Szendroi note, even a Russellian account could incorporate pragmatic considerations about evaluating sentences for truth, and thus could be made compatible with their results. Nonetheless, this first experimental exploration provides important initial insights and calls for further in-depth investigations.

The second set of studies I'd like to review here has not played much of a role, if any, in the theoretical discussions of definites in linguistics or philosophy, or, for that matter, even in psycholinguistics.[9] The first, well-known in the memory literature, were carried out by Elizabeth Loftus and colleagues. The primary interest was the effect of being exposed to linguistic expressions relating to a witnessed event on the memory representations that are formed about that event. One of the manipulations in doing so was to vary whether certain information was introduced by a definite or an indefinite. For example, in Experiment 4 of Loftus (1975), subjects were shown a three minute film clip of an accident involving a car running into a baby carriage. After seeing the film, three different groups of subjects had to answer different questions. For the first group, the five critical questions contained a definite presupposing the existence of some object that was not in fact part of the film (e.g., 'Did you see the children getting on the school bus?'). The second group instead saw a direct question with an indefinite about the presence of that object (e.g., 'Did you see a school bus in the film?'). The third group served as a control condition and only saw filler questions. One week later, the subjects came back in and now all had to answer the direct questions. The subjects in the definite question group incorrectly answered 'yes' significantly more often (29.2%) than both the direct question group (15.6%) and the control group (8.4%) (also see Loftus & Zanni, 1975;

[9]The studies were brought to my attention by the very accessible discussion of presuppositions and their role in real life communication in Sedivy & Carlson (2011).

Loftus et al., 1978). What these experiments demonstrate is that merely having been asked a question with a non-referring definite right after watching a film substantially increases the chances of having formed an incorrect memory representation of the film, as compared to having seen the same information presented with an indefinite. This fits naturally with a presuppositional view of definites, which accords different status to the same information in these two cases.

A more recent set of studies by Fiedler et al. (1996) follows up on a number of possible criticisms of the interpretation of these results by Loftus and her colleagues. Most relevantly for us, they argue that

> '[p]resupposing means to grant an aspect of information as a given if only for a transient memory representation, and this process is likely to turn on constructive processes. [...] This side effect [of building a representation of the presupposed entity; FS] of comprehending a presupposition is independent of its pragmatic success; it may even occur for correctly denied, false presuppositions' (Fiedler et al., 1996, p. 501).

Their crucial Experiment 3 employs an approach similar to the Loftus studies. Subjects are shown a video documenting the objects present in an apartment and subsequently are asked to evaluate a number of statements containing definites ('The shopping basket in the corridor was made of bast.') or indefinites ('A shopping basket was standing in the corridor.'). After this experimental treatment part of the task, data on the critical dependent measure was collected in the form of a recognition task, where subjects had to decide for each object on a list whether or not it was part of the apartment they saw. Different groups of subjects had to either make a simple truth-value judgment evaluating the statement relative to the film or to decide whether the statement was about the film they saw or another one (they were told that there are two such films). The second task variant bears on certain pragmatic explanations of the presupposition effect, which will be of some importance in the General Discussion section below. In a nutshell, subjects might be influenced by some broadly speaking Gricean reasoning in a way that leads them to assume the presupposition to be true if they hear it as part of a statement presented by the experimenter. But if the task is to decide whether the statement is about the movie they saw or another one, no such pragmatic effect should come into play.

The results reported by Fiedler et al. (1996) lend further support to Loftus's constructionist interpretation, where the encounter of information in a statement or question after visually perceiving a scene affected the memory subjects constructed for this scene. This effect is particularly pronounced when that information is presupposed.[10] Independent of the evaluation task, false recognition of objects (i.e., incorrectly stating that a non-object was part of the movie) was significantly more

[10]They did not find an effect of the control condition with the indefinite in Experiment 3, but did in other experiments, though it seems to generally be smaller than for presuppositions.

frequent for definite statements (32%) than indefinite ones (16%). Note that subjects were generally quite successful at recognizing objects that were actually present in the movie (with close to 90% accuracy), and that the correct rejection rates of the initial statements during the experimental treatment were quite high as well ($\approx$ 85%), regardless of the statement form and judgment type.

All in all, these studies provide intriguing evidence that definites involve a type of information that is not present with indefinites, and that this type of information is particularly apt at intruding processes of memory construction. A presuppositional account of definites offers a natural perspective on the effects found in these studies. If presuppositions consist of information that is treated as being taken for granted, and which is not the main point at issue, it would only seem natural that interpreting a statement containing presuppositions would make it particularly difficult to discard the information introduced by them subsequently. In Fiedler et al.'s (1996) words, 'presuppositions are efficient means of encouraging representations in which presupposed objects are taken for granted as given' (p. 507). Classical Russellian accounts, on the other hand, do not provide a distinction between definites and indefinites in terms of the types of content involved, and thus would need to be supplemented to explain these experimental findings. Furthermore, Fiedler et al.'s (1996) arguments against broadly speaking pragmatic accounts remain challenging for possible pragmatic supplementations of a Russellian account. Their relevant results also bear on the possibility of seeing the effects found in the studies reported below as reflexes of the process of local accommodation. We will turn to these points in more detail in the general discussion section. In the next section, we present two new experiments whose general approach shares some of the properties of the previous studies, but which crucially add a novel perspective by looking at response time data for truth-value judgments on sentences containing definite descriptions.

3 The time-course of truth-value judgments involving definites

3.1 Experiment 1

The first experiment, though actually carried out prior to the present author becoming aware of the experiments from the memory psychology literature reviewed above, takes a similar approach and provides further empirical evidence pertaining to the status of information conveyed by definite descriptions in contrast to indefinite ones. In particular, we set out to test whether rejecting a statement as false based on the non-existence of an individual matching the noun phrase description takes longer for definites than for indefinites. While this is generally in line with the memory effects reported above, it adds another dimension by looking at the initial process of rejecting a sentence based on a putative presupposition.

3.1.1 Design

The experimental design involved presenting subjects with a simple and concrete visual context and having them evaluate sentences with definites or indefinites that varied in status with respect to the crucial existence claim. In the critical conditions, the description did not pick out any of the entities presented in the visual context, whereas it did in the control conditions. Subjects had to make a forced choice truth-value judgment, with 'true' and 'false' as the only response options. We expected subjects to resort to the 'false' option in the critical condition even for the definite sentence version, based on the assumption that this could represent a general notion of rejection (see discussion and references in section 2.1).

The visual contexts consisted of arrays of colored shapes, and the sentences contained definite or indefinite noun phrases with a shape-label as a noun (e.g., 'circle', 'square', etc.). These were furthermore restricted by a post-nominal prepositional phrase specifying a location in the array (e.g., 'on the top'). The main predicate of the sentence always contained a color adjective (e.g., 'is green').[11] By using prepositional phrases with locations that can only be sensibly interpreted relative to the display at hand, we gained a tighter control over the context relative to which existence would be evaluated, which ensures that no potential referent instantiates the properties specified by the description. For the indefinite condition, we chose the existential 'there'-construction to ensure a purely existential interpretation and avoid a possible referential interpretation of the indefinite. An additional control using *exactly one*, to be discussed below, was also included. Two versions of the array of shapes varied the position of the critical shape. Example arrays and sentences are provided in Figure 1 and (5) respectively.

(5) **Examples of Target Sentences**

a. The circle on the left was green. **Def**

b. There was a circle on the left that was green. **Indef**

c. There was exactly one circle on the left and it was green. **Ex1**

Note that half of the items differed from the example used for illustration here in that they contained two circles, yielding variants where the triangle on the top right instead would have been a circle in a color other than green.

One worry in repeatedly presenting infelicitous sentences of the sort relevant here to subjects is that they notice the nature of the anomaly and adopt strategies in dealing with them that do not reflect normal circumstances. In order to avoid this, we attempted to make the experiment a bit more taxing and the task more

[11] This configuration was chosen to minimize potential pitfalls and confounds with respect to the existence requirement. For example, preliminary testing had indicated that using a color adjective to restrict the noun runs the risk of allowing an interpretation independent of the shape-array under consideration, e.g., in terms of a general type-level interpretation, where 'the green circle' refers to some abstract notion of a green circle.

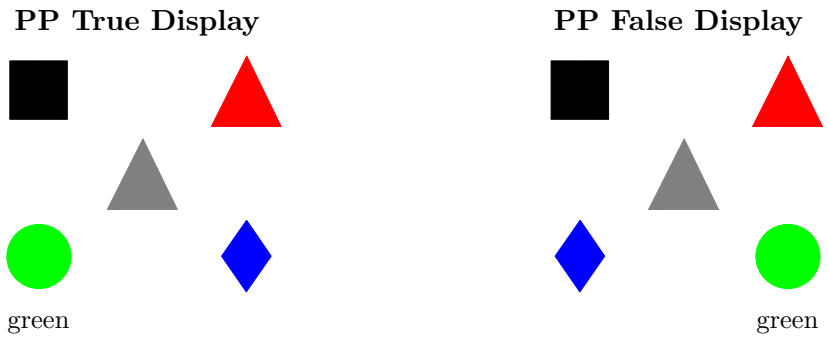

Figure 1: Example arrays

interesting by setting it up as a memory exercise. This introduced some degree of uncertainty about how the linguistic form relates to the display previously seen. The relatively demanding nature of the task also should keep subjects from reflecting on the status of the critical definite sentences, should they notice their anomaly, as their full attention is required continuously for the task at hand. The visual array was presented only briefly, followed by the presentation of the sentence to be evaluated (thus the choice of past tense in the sentences). This made the experimental task relatively demanding, but subjects performed reasonably well overall.[12]

Note that the color of the circle in both versions of the array did indeed match the adjective in the sentences. Thus, the decision on the truth-value judgment rested solely on the descriptive content of the noun phrase, specifically the locational prepositional phrase. When the latter matched the position of the relevant shape in the array, the sentence was clearly true. When it did not, it was clearly false for the indefinite sentence, and false or infelicitous for the definite version, depending on one's theoretical assumptions. Either way, it seems intuitively plausible to judge the sentence as 'false' within the present task.

While the indefinite control sentence is largely on par with the definite one in terms of the overall meaning conveyed, and specifically with respect to the existence requirement, there is one difference that both Russellian and presuppositional accounts would generally posit: they assume the definite to also introduce a uniqueness condition, which is not present for the indefinite. This could be relevant for verification strategies employed in the task and thus poses a potential problem for interpreting the data. To cover this possibility, the experiment utilized a third sentence condition, **Ex1**, that explicitly introduced uniqueness as part of the asserted content.

All three sentences were presented in the same two visual contexts, and therefore were matched in terms of the truth values of the expected responses. The central

[12]For the reader concerned about the memory component of the task, see experiment 2 below.

	Predictions		
Sentence/Array Type	**Def**	**Indef**	**Ex1**
(i) **PP true**	true	true	true
(ii) **PP false**	false (Russell)/ infelicitous (Presuppositional)	false	false

Table 1: Predicted Status of Sentences

comparison of interest is whether the relation between the two Array Types differs based on which Sentence Type we are considering. The controls introduce the existence (and uniqueness) requirement as part of the main asserted content. On a presuppositional account, the status of this requirement is different with definites. On a Russellian account, on the other hand, it also is part of the main asserted content. Russellian accounts thus predict the two array conditions to be in the same relationship to one another for all three Sentence Types. A presuppositional account, on the other hand, would lead us to expect an interaction of the array and Sentence Type factors, since the 'false'-judgments with definite sentences are based on presupposition failure, rather than plain falsity.[13] More specifically, if rejecting a sentence based on presupposition failure is harder than judging it false due to its truth-conditions not being met, response times for the former judgment should incur a greater delay relative to the control array (where the existence condition is met) than the latter.

To sum up the design, combining the sentences and arrays yielded an overall 2×3 interaction design with Array Type ((i) shape in location specified by PP (≈ **PP true**) vs. (ii) shape not in location specified by PP (≈ **PP false**)) and Sentence Type (definite (**Def**) vs. existential (**Indef**) vs. existential + uniqueness (**Ex1**)) as factors. The main comparisons of interest were the 2×2 interactions between Array Type and the definite sentences and the two control Sentence Types respectively. The theoretical status of the critical sentences is illustrated in Table 1.

3.1.2 Methods

Materials Using a set of six shapes (circles, crosses, diamonds, hearts, squares, and triangles) and six colors (black, blue, gray, green, red, and yellow), 36 items with versions in all six conditions described above were created and split into six lists containing six items per condition. This created a fully counter-balanced design, where each subject saw six items in each condition, and each item was seen in all conditions by different subjects. The location of the critical shapes was varied systematically and the location specified in the presuppositional phrases was adjusted accordingly. In addition to the 36 experimental items, each list contained 126 other items from

[13] It doesn't strictly speaking predict it, since we can't exclude the logical possibility that even if the underlying processes are different, the time course for them taking place could be the same.

other experiments. 36 of these were equivalent to the experimental items, except that the color adjective appearing as the main predicate was not true of the shape whose shape-label was used in the definite description.[14] Another 72 items involved sentences with definite plurals and arrays with nine occurrences of the same shape in different colors (discussed in Schwarz, 2013), and 18 items contained similar arrays followed by sentences containing the quantifiers 'few' or 'a few'. The order of presentation of items within each list was randomized for each subject, with no more than two subsequent trials from the same sub-experiment.

Procedure & Participants The experimental design was implemented using the Experiment Builder software package by SR Research and responses and response-time data were recorded.[15] Subjects were seated in front of a computer screen and received the initial instructions below, followed by a practice trial.

(6) **Instructions**

In this experiment, you will be shown simple pictures containing various shapes for a brief period of time. Afterwards, you will see a sentence, and your task is to evaluate whether the sentence is true or false relative to the displayed array. Try to push the appropriate button as quickly as possible.
We will begin with a brief practice trial. Then we have to set up the eye tracker. [See footnote 15]
If you have any questions, please feel free to ask the experimenter now!
When you are ready, press a button to proceed to the practice trial.

Each trial consisted of the following sequence of events:

(7) Sequence of events during each trial

a. Display of dot in center to control for initial eye position
b. Display of array of colored shapes for 900ms
c. 10ms pause
d. Display of dot in center to control for initial eye position
e. Display of sentence
f. Button press to indicate 'true'/'false' answer

[14] This sub-experiment did not yield significance for any of the key interactions. This is likely due to the fact that the relevant type of shape (and, in most cases, no other shape) was not of the color mentioned in the sentence. This allows for answering based on the color adjective alone, which likely is easier to remember than the combination of shape and location information. Since the color predicate was part of the asserted content, this would lead to no differences in response behavior across Sentence Types.

[15] We also recorded eye movement data using an EyeLink 1000 eye tracker, but the reading time results were not informative for current purposes; see footnote 18.

Responses were recorded using a ResponsePixx button box with 5 buttons. The left and right buttons were labeled as 'true' and 'false', with position of these values counter-balanced between subjects. After participating, subjects received a debriefing with a short explanation of the purpose of the experiment. 49 undergraduate students at the University of Pennsylvania participated for class credit.

3.1.3 Data analysis

All analyses used mixed-effect models with subjects and items as random effects, using the *lmer* function of the *lme4* package in *R* (Bates, 2005). Given recent arguments by Barr et al. (2013) that maximal random effect structures should be used whenever possible, we generally computed models with the maximal random effect structure that would converge, with random effect slopes for each factor, as well as the interaction where applicable. To assess whether inclusion of a given factor significantly improved the fit of the overall model, likelihood-ratio tests were performed that compared two minimally different models, one with the fixed effects factor in question and one without, while keeping the random effects structure identical (Barr et al., 2013). We report estimates, standard errors, and t-values for all models, as well as the χ^2 and p-value from the likelihood-ratio test for individual factors. For analyses of accuracy, which involved a binomial dependent variable, we report Wald's z and p-values from the *lmer*-output. To facilitate presentation of results, we will adopt the following shorthand indications for which random effect structure (RES) was used in the *lmer*-syntax in a given case:

- RES-1: $(1 + factor1 * factor2|\ subject) + (1 + factor1 * factor2|\ item)$ (Full model)
- RES-2: $(1 + factor1 * factor2|\ subject) + (1 + factor1 + factor2|\ item)$
- RES-3: $(1 + factor1 + factor2|\ subject) + (1 + factor1 + factor2|\ item)$

For 2×2 interaction analyses, predictors were centered, so as to render estimates of main effects. Comparisons between individual conditions were conducted using the appropriate treatment-coding, with reference levels as specified below for each comparison. I first present data treatment and accuracy rates and then move on to response and reading times.

Data Treatment & Accuracy Given the somewhat difficult nature of the task, which required memorization of five shapes and colors, accuracy levels were moderate, but sufficient for carrying out further analyses. The overall mean accuracy rate for items from the experiment reported here (where accuracy was coded according to the expected responses laid out in Table 1) was 75.3%. There was substantial variation in accuracy between subjects. Since only correct responses will be informative for looking at response time data, subjects with an overall accuracy rate of

	Def	**Indef**	**Ex1**
PP True	83.4 (76.0)	75.9 (71.9)	72.2 (65.8)
PP False	81.3 (77.0)	82.5 (77.1)	86.9 (84.9)

Table 2: Accuracy rates in % after exclusion of low-accuracy subjects. Accuracy for the full data set, before exclusion, shown for comparison in parentheses.

Interaction between Array Type and...	RES	β	SE	t	χ^2	p
...Sentence Type for **Def** vs. **Indef**	2	246.3	105.0	2.35	6.21	**<.05**
...Sentence Type for **Def** vs. **Ex1**	2	209.5	97.1	2.16	5.33	**<.05**
Post-hoc Comparisons (with Bonferroni correction):						
Simple Effect of Array for **Def**	2	-229.4	69.6	-3.30	5.68	**<.05**
Simple Effect of Array for **Indef**	2	15.6	89.1	0.17	0.03	>.8
Simple Effect of Array for **Ex1**	3	-9.1	75.0	-0.12	0.02	>.9

Table 3: 2×2 Interactions between Array Type and Sentence Type and Post-hoc Comparisons for Response Time Data in Expt 1

less than 66% were excluded from the data analysis, as their performance is not sufficiently above chance level. There were 10 such subjects, which left data from 39 subjects for purposes of data analysis. Furthermore, response time data points that were more than 2 standard deviations from the mean were also excluded from the response time data analysis. This affected 69 data points, or 4.9% of the remaining data.

The mean accuracy for experimental conditions is illustrated in Table 2, both before and after data removal based on accuracy. Accuracy was lowest for the **Ex1** condition when the PP matched the shape location (**PP true**). A logistic regression mixed effect analysis of accuracy (after exclusion of low-accuracy subjects) revealed this to be significantly different from the **PP False** condition ($z = 2.87$, $p < 0.05$ after Bonferroni correction). The difference between the Array Types was not significant for **Def** and **Indef**. This led to a significant interaction between Array Type and Determiner for **Def** for **Ex1** ($z = -2.64$, $p < .05$ after Bonferroni correction). There was no significant interaction between **Def** and **Indef**.

Response Times Response time was defined as the time that passed between the initial display of the sentence on the screen and the button press indicating the truth value judgment. Only trials with accurate responses were included in the response time analysis. Mean response times per condition are illustrated in Figure 1.

The central question was whether the nature of the switch from true to false

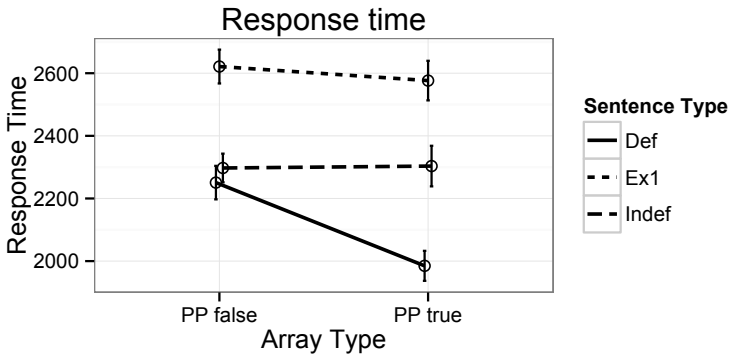

Figure 2: Response Times for Correct Responses

in the **Def** condition was different from that in the **Indef** and **Ex1** condition. Statistically, this corresponds to asking whether there were interactions between Sentence Type and Array Type when comparing **Def** to the other two conditions.[16] A mixed-effect model regression analysis revealed both interactions in question to be significant, as detailed in the upper part of Table 3. To further investigate the nature of these interactions, post-hoc comparisons between Array Types were conducted for each Sentence Type, by using treatment coding with **PP False** and the corresponding Sentence Type set as reference levels.[17] Only **Def** showed a significant difference, with faster response times in the **PP True** condition, as shown in the lower part of Tabel 3.[18]

3.1.4 Discussion

The present experiment was designed to test whether reaching a false judgment involves a different process when it is based on the falsity of the existence condition

[16]Main effects were not analyzed. The effect of Array Type is dominated by the interaction, and the differences between the sentence types are expected and theoretically uninteresting, given the differences in sentence length.

[17]For example, to test the effect of Array Type in the **Def** conditions, the **Def** level of the Sentence Type factor was coded as 0, which renders the Array Type coefficient in the *lmer*-output to express the impact of Sentence Type at the **Def**-level.

[18] While the response times were the primary focus of the present study, we also inspected eye tracking data on reading times to test for early reflexes of the effects. Standard reading time measures for the noun phrase region (e.g., 'circle on the left') and the position adjective on its own (e.g., 'left') were computed. While there were significant differences based on Sentence Type (with faster total reading times in the **Ex1**- and **Indef**-conditions) - attributable to the varying sentential contexts -, none of the effects discussed above for response times (nor any other effects) were significant, apart from a marginally significant slow-down on *left* in the **Def-PP false** condition (consistent with the simple effect of Array Type for **Def** in response times).

of a definite. We addressed this question by looking at response times, comparing definites to indefinites and true judgments to false judgments. A Russellian analysis assumes that existence is part of what is asserted in both cases, whereas presuppositional accounts assume it to be presupposed in the case of the definite. While the type of response given for the two cases was on par (namely 'false'), response times were longer for non-referring definites in comparison to the Array control conditions relative to the same comparison for indefinites. Before interpreting this result as directly bearing on the presuppositional status of the existence component of definite descriptions, we need to consider a number of potential issues with this experiment that might undermine such an interpretation.

One possibly worrisome observation is that the difference found for **Def** in the different Array Types could simply be attributed to a standard response bias effect: it is well known that 'false' responses for (affirmative) sentences generally take longer than 'true' responses (Clark & Chase, 1972; Gilbert et al., 1990; Gilbert, 1991). The remarkable outcome then would be that no such response bias is found in the **Indef** and **Ex1** conditions. To maintain an interpretation that attributes the interaction to the presuppositionality of the definite, we'll need to assume that the response bias is underlyingly present in both types of conditions, but that it is masked by another factor counter-acting it. And there indeed is a highly plausible candidate for such a factor, namely the timing of information crucial for falsification and verification: in the **PP False** conditions, it becomes clear early on, upon reading the PP '. . . on the left', that the sentence is false. In contrast, the **PP True** condition requires reading the color adjective at the end of the sentence for judging the sentence 'true'. If subjects initiate the judgment process in the former case right upon encountering the PP, this would lead to a decrease in response time in the **PP False** condition, relative to the **PP True** condition. This relative speed-up does not show up in the response times for **Indef** and **Ex1** because it is counter-acted by the standard response bias, with a relative speed-up for the 'true' responses in **PP True**. Both factors presumably are present for **Def** as well, i.e., the early introduction of false information in the **PP False** condition should introduce a speed-up that cancels out the standard response bias. But if that were all that was going on, we should get a result entirely parallel to **Indef**. The fact that **Def** nonetheless exhibits a delay for **PP False** then must be due to another factor, arguably the presuppositional status of the information inside of the PP. While this interpretation of the data seems highly plausible on its own, Experiment 2 below will provide further support for this view.

A second potentially problematic aspect of the results is that we do not actually find a direct difference between 'false'-judgments in the **Def** and **Indef** conditions, but only have a relative difference in comparison to the **PP True** conditions, as reflected in the interaction. Note, however, that the lack of a difference in the **PP False** conditions has to be taken with a grain of salt, as the response times include the time for reading the sentences, and the **Indef** sentences were systematically

longer than the **Def** sentences. What we ideally would want is a way of lining up the times at which the falsifying information becomes available in each Sentence Type, to provide a fair comparison. Experiment 2 allows for exactly that, as the linguistic stimuli are presented auditorily and response times can be measured directly from the point in the stimulus where the falsifying information becomes available.

Yet another potential issue concerns possible effects of contrastive inferences for items such as the one illustrated in (5): It is well-known from the literature, largely based on work within the visual-world paradigm (e.g. Crain & Steedman, 1985; Sedivy, 2003; Heller et al., 2008) that modified definites give rise to contrastive inferences - the expression *the circle on the left* might thus lead to the inference that there is another circle in the display. And if there only is one circle present, this could lead to delays in some way or other based on the fact that this inference is not matched in the display. Since half of the items actually did contain a second circle, we are in a position to test the effect of this contrastive inference straightforwardly by including a between-item factor encoding whether there was one or two shapes of the relevant sort. While the effect appeared to be somewhat stronger numerically for the items where the contrast-issue was present, a 3-way interaction analysis including Sentence Type, Array Type, and the contrast factor did not yield a significant overall interaction. Moreover, the 2-way interactions of Sentence Type and Array Type remained essentially unchanged. It thus does not seem to be the case that the crucial 2×2 interactions can be attributed to this issue of contrastive inferences that were not matched in the display. Furthermore, the items in experiment 2 all had two entities of the relevant type present in the displays used there, so that the potential issue of contrastive inferences did not arise at all.

A final point of potential worry concerns the memory component of the task and the resulting level of difficulty, which led to the exclusion of a fair amount of data. While it is not necessarily clear how the mere fact that we used a somewhat taxing task involving memory would lead to the response time pattern that we found, it nonetheless would indeed be reassuring to find comparable results in a simple task without a memory component and with higher overall accuracy. Experiment 2 addresses this issue as well.

3.2 Experiment 2

3.2.1 Design

While the general spirit of the design of Experiment 2 was entirely parallel to that of Experiment 1, a number of changes were implemented to address the potential worries mentioned above. Rather than having a sequence of visual and linguistic stimuli presented on a screen, a visual display was presented in parallel with auditory linguistic stimuli. Thus, there was no memory component involved in the task. Furthermore, the auditory stimulus presentation provided a direct way of pin-pointing the point in time when the falsifying information was introduced, thus

making it possible to calculate response times relative to that. The design was also adjusted in such a way that all conditions required a 'false' response, thus factoring out any potential issues related to response bias. In combination with the more fine-tuned response time calculation, this also allowed for a meaningful comparison of how 'false'-judgments arise in the definite and indefinite conditions. Finally, the inclusion of true control conditions, which were primarily intended to serve counter-balancing purposes, allowed a direct assessment of the presence of response-bias effects in the indefinite conditions.

3.2.2 Methods

Materials The materials for the second experiment were of a slightly different nature, in that they did not involve abstract, colored shapes, but rather contained pictures of people together with a calendar strip that included iconic representations of activities. These were introduced as reflecting the plans of the relevant people for the coming week. This made for a more naturalistic set of materials, while still allowing for tight control over the properties associated with each individual and ensuring that the sentences were evaluated only with respect to the current display. Each display contained four individuals, with two each of two different types of people (boys vs. girls and men vs. women), as illustrated in Figure 3 (with the two possible variants of the target picture discussed below). The definite sentences (**Def**) had the form 'The X with an outing on Y is going to Z', where 'X' was a noun for one of the four categories of people, 'Y' was a weekday, and 'Z' an activity or location of some sort. The indefinite sentences (**Indef**) minimally differed from that in the same way as before. One of the pictures was manipulated in such a way that either the weekday or the activity did not match the sentence, thus rendering it false based on either the prepositional phrase or the main predicate of the sentence.[19] The other picture with a person of the same category served as a distractor, as did the two other pictures of people of a different type. The three distractors always involved different weekdays than either version of the target picture to avoid any confusion. 24 items with versions in each of the four conditions were created. Example sentences from the materials is provided for illustration in (8), which were paired with the visual stimuli in Figure 3:

[19]Note that it is important to make just one of these two false in the respective cases. The sub-experiment within Experiment 1, mentioned in Footnote 8, likely did not yield any relevant significant results because the variation was between sentences that were either false due to asserted content alone or due to both asserted and presupposed content. Rejections of the latter cases thus could have been due to either one of the two false components.

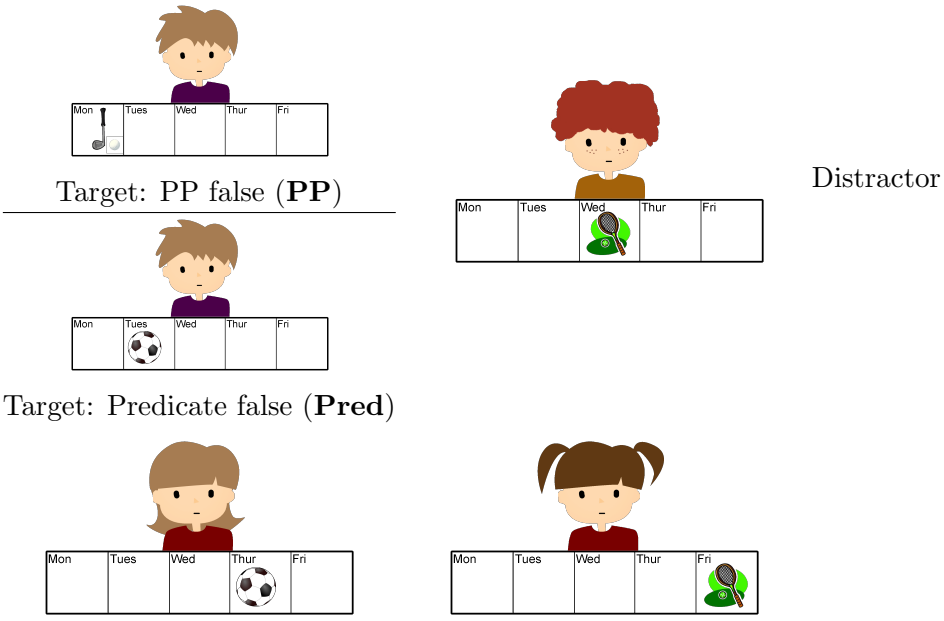

Figure 3: Sample Display for Experiment 2 showing both possible variants of the target picture

(8) a. The boy with an outing on Tuesday is going to play golf. (**Def**)
b. There's a boy with an outing on Tuesday who's going to play golf. (**Indef**)

The days of the week were varied across items, as was the positioning of the pictures, with each pair of people of the same category horizontally aligned. The sentences were recorded and prepared in Praat by a research assistant. The same recordings of the sentences were used for both display types. There was a total of 120 items in the overall experiment. In addition to the 24 critical items, there were 24 equivalent items as controls, which involved sentences that accurately described one of the people in the display. Furthermore, there were items with similar displays, but containing activities on multiple days, and sentences with the quantifiers *a few*, *only a few*, and *not many*. Half of these were false, whereas the other half was either true or false, based on whether a literal or non-literal meaning of the quantifier was adopted. Overall, this rendered a slight bias towards 'false' answers in the experiment as a whole. Subjects saw one of four lists of materials containing one version of each item, with a fully counter-balanced set of data points to be collected for each item and condition. The order of presentation of items was randomized for each subject, with no more than two subsequent trials from the same condition.

Procedure & Participants The experimental design was again implemented using Experiment Builder, though only responses and response time data were recorded. Subjects were seated in front of a computer screen and provided with headphones. They then saw the instructions below:

(9) **Instructions**

In this experiment, you will see a number of pictures displayed on the screen. You then will hear a sentence describing parts of the display and your task is to decide whether or not this description indeed holds of the display.

If it does, press the button labeled 'True', and if not, press the button labeled 'False'. Please try to make your decision and press the button as quickly as possible.

To continue and do a practice trial, press any key on the keyboard. If you have any questions, feel free to ask the experimenter after the practice trial.

[practice trial]

Thanks! As we continue, don't forget to examine all of the pictures when determining the truth of the sentence. Also, please try to give your answer as quickly as possible.

Let's continue the experiment!

The presentation of each trial began with the presentation of the visual display. After 500 ms, playback of the recording of the sentence was initiated, and subjects had to push a button to enter their response (it was possible to respond prior to the end of the recording). Responses were recorded using either a ResponsePixx button box or a DirectIN keyboard, both of which allow for millisecond accuracy in response time recording. The left-right distribution of 'true' vs. 'false' answers was varied across subjects.[20] After participating, subjects received a debriefing explaining the purpose of the experiment to them. 28 members of the University of Pennsylvania participated in the experiment, mostly for class credit, and in a few cases on a voluntary basis.

3.2.3 Data analysis

The same types of statistical analyses as described above were carried out for Experiment 2. Accuracy of subjects' responses was consistently high (above 95%) and did not differ across conditions. Response times were calculated relative to the onset of the critical part of the sentence. While subjects generally did not respond prior to the end of the sentence, they did so a fair number of times in the **PP** conditions (responses prior to the onset of the sentence-final activity or location in the

[20] Due to a logistical error, this ended up not being entirely counter-balanced, with 19 subjects having 'true' responses on the left, and 9 on the right. However, since our critical comparison only looks at 'false' responses, this should not skew the data in any way.

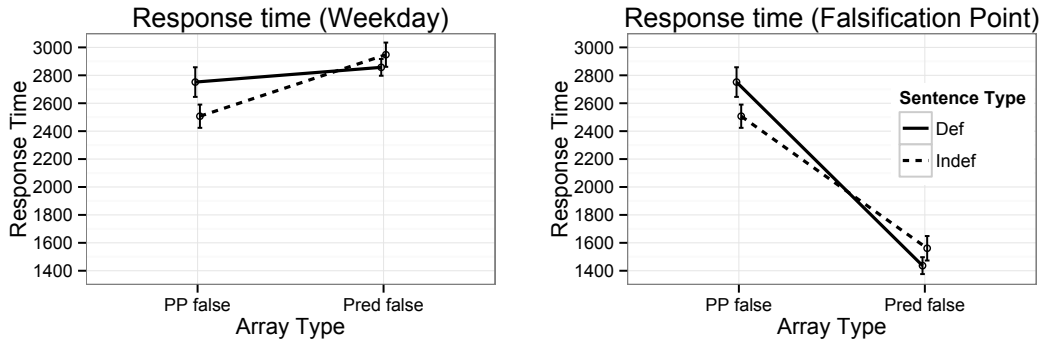

Figure 4: Response Times for correct responses (measured from the onset of the weekday mentioned in the sentence and from the point of falsification respectively)

Pred condition were excluded, since they were made prior to the crucial information). Only response times from trials with correct responses were included in the analysis.

Key Interaction As a first approximation for purposes of illustration, the mean response times as measured from the onset of the weekday are presented in the left panel of Figure 4. But note that this does not provide an ideal comparison across Sentence Types, as the falsifying information relative to the picture is introduced at different points in the sentence, as detailed in Table 4. Furthermore, the average duration from the onset of the weekday to the onset of the last word was 31 ms longer in the **Def** recordings, which means that the value for the **Indef-Pred** condition is under-estimated by this amount (relative to **Def-Pred**) when measuring from the weekday onset. For the purposes of statistical analysis, we therefore used the more appropriate measures based on the point of falsification in each condition. The mean response times based on this measure are graphed in the right panel of Figure 4. Note that the response times for the **PP** conditions are substantially higher here relative to those in the **Pred** conditions because they were calculated from an earlier point of reference.

As in Experiment 1, our main question of interest is whether there is an interaction, which reflects whether the basis of the judgment (falsity of PP or main predicate) affects response times differently for definites and indefinites. A mixed effect-model regression analysis indeed reveals a significant interaction between the two factors, as detailed in Table 5. Thus, the magnitude of the change between the **Def** and **Indef** conditions differed significantly based on whether we're looking at the **Pred** or **PP** conditions, as predicted on a presuppositional account. The

	Falsification Point	
	PP	**Pred**
The boy with an outing on	Tuesday is going to play	golf.
Def	1288	2707
There's a boy with an outing on	Tuesday who's going to play	golf.
Indef	1442	2830

Table 4: Mean onsets (in ms) of falsification points by condition

	RES	β	SE	t	χ^2	p
Sentence×Array Type Interaction	1	361.39	170.56	2.12	4.29	**<.05**
Post-hoc Comparison:						
Def vs. **Indef** (**PP** condition)	2	238.9	107.3	2.27	4.78	**<.05**

Table 5: Interactions and Effect of Sentence Type for **PP** in Expt 2.

presuppositional account also predicted **PP** response times in the **Def** condition to be slower than in the **Indef** condition. This prediction is borne out as well, as revealed by a Post-hoc comparison using treatment coding with **Indef** and **PP** as reference levels (see second line of Table 5 for details).

The present results complement those from Experiment 1 rather nicely. By controlling for the timing of the introduction of information that allows falsification of the sentence, we now have a better comparison between the **Def** and **Indef** conditions, and indeed find a significant delay for 'false' responses for the former. Crucially, there seems to be a clear effect of introducing the relevant information earlier, as shown by the difference in the **Indef** conditions, which in the present case counter-acts the delay for 'false' response based on presupposition failure in the **Def** condition. Put most simply, the results in the present experiment are equivalent to those in Experiment 1 if we simply add a constant to the 'true' conditions there to account for the standard response bias effect (and shift the **Indef** times up since response times are increased due to the longer sentences).

Additional Follow-up Analyses To further substantiate this perspective on the data, we used the true fillers from the present experiment to confirm that there indeed is a response bias effect for the indefinite sentences. Looking at the subset of the data containing only the **Indef** conditions (with displays of the **PP**, **Pred**, and **True** kinds), we find a significant difference between **Pred** and **True** ($\beta = 335.31$, $SE = 105.96$, $t = 3.16$), with longer response times in the former. The mean response times as measured from the onset of the last word of the sentence for both

	Pred	PP	True Fillers
Def	1437	1332	1224
Indef	1561	1121	1227

Table 6: Response times from the time of the onset of the last word

critical conditions and true fillers are detailed in Table 6. Note that response times in the **Indef-PP** condition are numerically smaller than those in the **True** condition, despite being judged 'false'. They thus exhibit exactly the type of speed-up that is expected if the falsification process is initiated earlier, based on the time at which the relevant information is introduced.

Finally, looking at the **True** filler and **PP** conditions for both **Def** and **Indef** together with a 2×2 interaction analysis (with Display Type as a between item factor) provides a replication of sorts of Experiment 1. While the interaction does not seem to reach full significance ($\beta = 235.05$, $SE = 136.63$, $t = 1.72$), this is likely due to a lack of power given the necessary between-item comparison. All in all, the results from the two experiments are very much consistent with one another and support the assumption that response times decrease when the information relevant for a 'false'-response is introduced early on in the sentence. This explains the absence of a response bias effect for **Indef** in Experiment 1, as well as the fact that there is no significant Array type effect for **Def** in Experiment 2 (when measuring from the same point, as in the left panel of Figure 4). In both cases, the effect of interest ends up being counteracted by the early introduction of the relevant information.

3.2.4 Discussion

Experiment 2 was designed to further substantiate the proposed interpretation of Experiment 1 by addressing a number of possible concerns that might undermine that interpretation. The design no longer involved a memory component, and removed the issue of response bias from the data by only looking at 'false' responses in all four conditions. The timing information was further refined by locking response times to the onset of the critical information in the auditory linguistic stimuli. Since there always were two individuals of the relevant sort in the display, we also did not run into the potential issue of contrastive inferences at odds with the displays. The results were entirely parallel to those from Experiment 1, once we put a crucial assumption into place: the assumption is that early introduction of falsifying information (here, in the PP) leads to a speed-up in response times. The validity of this assumption was supported by follow-up analyses. Furthermore, Experiment 2 provided direct evidence in terms of comparing 'false' responses based on presupposed vs. asserted content, with a corresponding delay in the **Def-PP** condition compared to the **Indef-PP** condition. Finally, we found direct evidence for re-

sponse bias effects in the **Indef** condition, addressing the potential concern about the apparent absence of such effects in Experiment 1, and at least have suggestive statistical evidence for a replication of that experiment by doing an analogous post-hoc between-item analysis that includes the true filler items. The results of Experiment 2 thus are consistent with the proposed interpretation of Experiment 1, and indeed provide further evidence for such an interpretation. Responding with 'false' takes more time when that response is based on presupposed information being false than when it is based on false asserted information.

4 General Discussion and Theoretical Implications

If we compare a classical Russellian account with one positing a conventionally encoded presupposition in light of the present results, the former clearly faces greater challenges. The latter can capture the differences in response times by claiming that rejection based on the falsity of presupposed content is more involved than rejection of asserted content. This goes naturally with seeing presuppositions as having a different status, e.g., by virtue of being backgrounded information or being taken for granted. While the notion that definite descriptions involve some type of presupposition may not be all that controversial amongst linguistic semanticists, the Russellian view is still commonly seen as the main competing account (for a recent example, see Elbourne, 2013, p. 5). One key argument in favor of presuppositional accounts is based on the questionable status of sentences with non-referring definites, which are claimed to not generally be straightforwardly false. However, it has turned out to be rather difficult to replicate the relevant intuitions systematically in a judgment paradigm, in particular with non-negated sentences. The present results based on response times constitute a novel approach to investigating the status of such sentences that doesn't rely on speakers' conscious decisions. They provide evidence for differences between presupposed and asserted existence requirements even in the absence of clear differences in the ultimate judgment.

Beyond the basic perspectives in the tradition of Russell, Frege, and Strawson considered so far, there are of course many modern accounts of definites. In relating our findings to the broader literature there are a number of avenues to consider. First, we'll explore what the source of the judgment delays (as well as the memory effects discussed earlier) might be from the perspective of conventional presuppositional accounts. This will directly relate to the literature on variation in truth-value judgments. Secondly, we'll consider possible amendments of a Russellian view, in particular in light of various neo-Russellian accounts that integrate a presuppositional component in one way or another.

4.1 Conventional presuppositions and the nature of the delay

Given the interpretation of the reaction time delay as a reflex of the presuppositionality of definites, any account based on a lexical entry for the definite article that encodes an existence presupposition would seem straightforwardly compatible with our results. This includes familiarity accounts in the tradition of Heim (1982), in particular ones that are based on weak familiarity (Roberts, 2003), which is essentially equivalent to an existence requirement. However, it is worth considering in some more detail just what the process leading to a delay is. This also relates directly to the question of how to account for plain 'false' judgments within a presuppositional account. I will consider three theoretical options for understanding the delay, namely in terms of squeamishness, special verification strategies, and (ultimate) interpretations that lack the presupposition.

Squeamishness Given the force-choice nature of the task - with 'true' and 'false' as the only options, our results leave open whether subjects would have given a 'false' response spontaneously, or whether they might have preferred to give a 'can't say'-type of response. In light of Abrusán & Szendrői's (2013) results for affirmative versions of their sentences, it is at least possible that 'false'-answers might be closer to the norm in reaction to presupposition failure than one might have previously thought. Either way, the delay could be due to subjects experiencing some squeamishness, either consciously or sub-consciously. For example, a Strawson-Reinhart-style account of truth-value judgments based on topicality would seem to be a good candidate for such an explanation (our definites appear as subjects in their sentences and could be seen as topical): on this view, subjects choose the 'false' response reluctantly - and after some consideration -, as a proxy for rejection on more general grounds. Alternatively, some form of squeamishness might be at play at an unconscious level, i.e., subjects might not necessarily be aware of any hesitation to judge the sentences as 'false', but such hesitation might be present unconsciously and slow down the steps involved at reaching a judgment. Schoubye's (2010) account based on questions under discussion fits naturally with this view, assuming that the question under discussion in the experimental task is (roughly) 'What were the shapes like in the array you just saw?'. This would make the statement a 'consonant conversational move' in Schoubye's terms and thus eligible for a plain 'false'-judgment.

Verification Strategies A second perspective on the data is that we are dealing with genuine 'false' judgments despite the presence of a presupposition, with no squeamishness involved. However, the 'false' judgment could still be arrived at in a manner different from rejections based on assertion. Proponents of accounts based on the assumption that such 'false'-judgments are due to the availability of some type of verification procedure (e.g. Lasersohn, 1993; von Fintel, 2004) could argue that verification with non-referring definites is in some way more involved than when

the existence condition is asserted.[21] In particular, the process of (what amounts to) accommodating the presupposition for the sake of going through the verification steps could plausibly be held responsible for the extra processing cost.

No Presupposition The final possibility to be considered here also sees the 'false'-judgments as genuine, but under the assumption that in such cases the interpretation of the definite does not involve a presupposition at all.[22] In order to explain the delay, such an account would have to assume that the presupposition-less interpretation is not arrived at immediately. Most plausibly, one could appeal to 'local accommodation' (Heim, 1983), where presuppositions are interpreted relative to operators whose scope they normally escape. And indeed, there are experimental results suggesting that local accommodation interpretations of verbs like *realize* and *stop* come with response time delays (see Chemla & Bott, 2013; Romoli & Schwarz, 2015, respectively). On this view, the delayed responses in our experiment could be seen as a confirmation of these findings. However, note that this would require the assumption that a simple 'false' response in effect is equivalent to asserting the negation of the sentence that is being responded to. While there are plausible proposals for analyzing response particles such as *yes* and *no* as involving ellipsis (e.g. Kramer & Rawlins, 2011), I'm not aware of any extension of such proposals to a response expressed as *false*. Thus, this option has to be spelled out in more detail in order to assess its theoretical viability.

On the empirical side, the results from the memory experiments by Fiedler et al. (1996) should also be considered in this regard. In the movie-selection version of their task, they found that incorrect memory representations based on the presupposition of a definite were frequently constructed even when the initial response to the relevant sentence correctly categorized it as pertaining to another movie. But it's unclear how local accommodation would come into play with this type of task, where there is no 'false' judgment to begin with. And even if there was an implicit step of judging the statement as 'false' with a locally accommodated presupposition, this should have an effect opposite of that observed, in that it would more firmly establish the falsity of the existence claim in question. Barring an alternative explanation of these results consistent with local accommodation, the previous perspectives would seem to have a slight advantage over a local accommodation account.

In sum, our results suggest that rejecting statements based on the existence requirement of definites is hard. This fits straightforwardly with presuppositional accounts of the type considered here: by its very nature of being presuppositional,

[21] But note that such accounts are also consistent with attributing the delay to squeamishness if they don't see the cases in question as involving genuine 'false'-judgments.

[22] Strawson argued along these lines, of course, but note that his proposal for taking topicality as the crucial factor for the absence of squeamishness presumably doesn't predict the sentences in our study to be non-presuppositional.

it is categorized as something that can and should be taken for granted. As Fiedler et al. (1996) put it, it necessarily has to be processed and at least temporarily adopted as true for the sake of simply interpreting the overall utterance it appears in. Therefore, it is unlikely to be critically evaluated itself (at least initially), and challenging its accuracy and detecting its falsity comes with additional processing efforts. And based on the results from the memory literature, the chance of it leaving traces that can affect memory formation even upon successful rejection of the presupposed information is higher than with literal, truth-conditional content. Various lines of attack are open to presuppositional accounts to give a more detailed characterization of the nature of the response time delay. Teasing these options apart will be an important next step for experimental work in this area.

4.2 Variations of Russellian accounts

So far, we have only considered a classical version of a Russellian account and concluded that it provides no direct explanation for the observed response time delays. However, there have been many modern variations of this account, which we briefly turn to now. For present purposes, we can group these together in three classes, which I will discuss in turn:

Referential Pragmatics The first variation of a Russellian account does not lend itself to explaining our results. Kripke (1977) argued that the referential interpretations of definites discussed by Donnellan (1966) can be accounted for in terms of a pragmatic supplementation of a Russellian account. But such referential interpretations, where the referent does not actually meet the noun phrase description, provide no help in explaining the present data. For example, in Experiment 1, subjects should have responded 'true' in the critical condition, if they interpreted *the circle on the left* to refer to the circle in the display (but not on the left), since it always had the color expressed by the predicative adjective. But the levels of 'false' answers were just as high in the **Def** condition as in the others.

Russell + Pragmatic Presupposition Another common modern adaptation of a Russellian account is to assume the truth-conditional contribution of definites to be as proposed by Russell, but to furthermore assume that there is a pragmatic presupposition of existence (and possibly uniqueness). Early characterizations of such a position can be found in Grice (1981) and Stalnaker (1974), and given the recent resurrection of pragmatic accounts of presupposition more generally (e.g. Simons, 2001; Schlenker, 2008, 2009), consideration of this possibility is all the more relevant.[23] The general idea would be that the perceived presuppositional

[23] As an anonymous reviewer points out, work by Sudo (2012) and Klinedinst (2012) also is relevant to the question of whether definite sentences entail existence, perhaps in addition to presupposing it.

effect standardly associated with the use of a definite can be derived on the basis of general principles of conversation (Grice, 1975).

The experimental results, both from the response time studies presented here and from the earlier memory literature, are difficult to account for on such a view. In particular, the sped up response times for the **Indef-PP** condition in Experiment 2 show that falsification based on asserted content begins right when the relevant expression is encountered. The question then arises why the asserted existence condition of the definite that is assumed by the accounts in question is not treated in the same way. That is, if the plain literal, truth-conditional content of an expression suffices for falsification, then why not utilize the same type of information in the same way in both the **Def** and **Indef** conditions, especially when providing a truth-value judgment is the sole task at hand?[24] The only possibility for explaining the delay from this perspective would be to claim that an optional pragmatic inference is computed and leads to a delay anyway. But it is entirely unclear why such a pragmatically generated presupposition should be considered in the experimental task, since the response can already be determined based on the literal content alone. So the challenge for proponents of a pragmatic account is to explain just why the pragmatic inference should come into play at all, and why the response based on the literal meaning doesn't proceed on the same time-line as in the case of indefinites.

Furthermore, the results from the memory literature are once again highly relevant. Fiedler et al.'s (1996) movie-selection task was designed to address the possibility that the memory effects are due to a type of pragmatic accommodation: when subjects hear the experimenter use an expression that pragmatically presupposes the existence of a shopping basket in the apartment in question as part of a truth-value judgment task, then subjects might be inclined to believe the presupposed proposition, due to his general authority and level of knowledge. However, in Fiedler et al.'s Experiment 3 subjects had to judge whether the sentences in question pertained to the apartment that they saw a movie about, or about another one. If they correctly concluded that the sentence was about another apartment, then they would have no grounds for pragmatically inferring the existence of such a thing in 'their' apartment - the conditions for triggering the pragmatic derivation of the presupposition as applied to the subject's apartment simply aren't met in such a case. Furthermore, the decision on which movie the sentence was about can be made based on the literal meaning alone, so it's unclear why any such inference should be considered at all. Nonetheless, even in cases where subjects initially classified the sentence correctly, they were more likely to incorrectly identify a shopping basket as something they had seen in their apartment a week later. This type of memory

[24] As a reviewer points out, the results from the other sub-experiment mentioned in footnote 9 suggest that this is indeed the case, as the false color adjective there seems to erase the effect of unmet presuppositions. However, since the results from that study are not significant, and a three-way interaction analysis between the relevant two sub-experiments did not reach significance, this is suggestive evidence at best.

intrusion effect posits another serious challenge for a pragmatic account.[25]

Existence + Familiarity Presupposition The final class of theories to consider here are ones that combine Russellian truth-conditions with a familiarity (or givenness) presupposition Szabó (2000); Ludlow & Segal (2004); Abbott (2008). At least on one construal of these theories (advanced explicitly by Abbott), the presupposition is not conventional, but rather of a pragmatic nature based on Gricean inferencing. Such a version of this type of theory faces the very same issue as other pragmatic accounts of presuppositions.

On another construal of this type of theory, the familiarity presupposition is conventional (this is suggested by Szabó, 2003, and may be compatible with Ludlow and Segal's proposal).[26] On the face of it, such an account still faces the challenge of why the asserted content of a definite is not utilized in the same manner as that of an indefinite, as can be seen from the different response time results in our studies. However, since it assumes the relevant presupposition to be part of the lexically encoded content of the definite article, consideration of the presupposition could be seen as compulsory and not dependent on suitable pragmatic circumstances. While falsification of a definite sentence based on the literal content might then be possible early on, the detection of the unmet presupposition could lead to delays in the response process, in ways similar to the options discussed above for other accounts based on conventionally encoded presuppositions.

Conclusion In sum, I submit that the experimental data under consideration does not lend itself to being explained by pragmatically supplemented neo-Russellian accounts. However, a neo-Russellian account that posits a conventionally encoded presupposition is not necessarily inconsistent with it. The upshot of our results, from this perspective, then, is that they lend support to the view that the existence presupposition of the definite article is conventionally encoded in its lexical entry. Beyond theories of definites, this is of obvious importance and relevance to current debates about the nature of presupposed content, where the idea of purely pragmatic accounts of presuppositions has seen a revival of sorts (e.g. Simons, 2001; Schlenker, 2008, 2009). The present discussion suggests that our results provide evidence for the existence of conventionally encoded presuppositions in the case of definites. Furthermore, the methodology introduced here opens the door for investigating the same question for other presupposition triggers as well.

[25] As an anonymous reviewer points out, additional options for a pragmatic approach open up if we assume the existence inference to be obligatory. For example, subjects could be aware of the definite's signaling of such an obligatory inference right away and be compelled to process this content more deeply, leading to a delay. While I can't rule out this option based on the present data, it seems counter-intuitive to me that special, deeper processing is applied to the presupposed content, based on the backgroundedness of presuppositions.

[26] While differing with respect to the nature of the presupposition(s) involved, the recent proposal by Hawthorne & Manley (2012) also would seem to fall into this general category.

4.3 Conclusion & outlook

To conclude, this paper has argued that response time measures can provide insights into the nature of the processes involved in assessing sentences with definites. This, in turn, allows us to confirm the existence of underlying distinctions of some theoretical interest. In particular, we were concerned with the difference between plain falsity and presupposition failure posited by presuppositional accounts. We found empirical evidence correlated with this difference, which supports an account of definites that assumes a conventionally encoded existence presupposition. While some circles may already subscribe to such a view quite uniformly, variants of Russellian accounts continue to have a strong standing amongst philosophers of language. Our results speak at least against the most common variant of such accounts that appeal to a pragmatic notion of presupposition.

Empirically speaking, the current approach using response times makes a methodological contribution as well, as there is little systematic empirical evidence to date substantiating the intuitions that presuppositional accounts are based on. In fact, the work of Abrusán & Szendrői (2013) shows that it is rather difficult to detect the relevant effects by investigating judgments alone (at least for affirmative sentences).

The results also bear on the debate in the literature on variation in truth-value judgments for definite sentences and provide first insights into the processes leading to rejection based on presupposed vs. asserted information. They open up the possibility of further investigating whether cases that are commonly considered to be false, rather than infelicitous, in the literature involve different processes that lead to only nominally equivalent judgments. In future work, consideration of various embedding environments, such as negation and conditionals, will also be of great theoretical interest, though additional issues arise once presupposition projection is involved (see the discussion of local accommodation above, and Schwarz & Tiemann, 2012, 2013; Chemla & Bott, 2013, for further details). Finally, it would be interesting to see whether the other standardly assumed presuppositional component of definites, namely uniqueness, behaves in ways similar to existence, or whether there is any difference between these (as suggested, for example, by Abbott, 2008). While the comparison is not straightforward, as violations of uniqueness can be remedied by domain restriction, initial results suggest that similar delays can be found for responses based on uniqueness violations (Schwarz, 2012).

Author's address

Florian Schwarz
University of Pennsylvania
Department of Linguistics
619 Williams Hall
255 S. 36th Street

Philadelphia, PA 19104-6305
E-mail: florians@ling.upenn.edu

Acknowledgments

This project has materialized over the course of a number of years, and I'd like to thank various people for helpful discussions along the way. Lyn Frazier and Angelika Kratzer provided input on early ideas for the project. Chuck Clifton gave extremely helpful advice on the interpretation of the data on various occasions. Shravan Vasishth provided helpful feedback on the statistical analysis. Participants of lab meetings in my lab at Penn and the Trueswell-Gleitmann lab, and of a seminar on experimental work on presuppositions at Penn in the fall of 2013, as well as audiences at a workshop at Rutgers in 2011 and at Linguistic Evidence 2014 provided stimulating discussion. Finally, several anonymous reviewers provided detailed comments on all aspects of the paper, which greatly improved its overall content and form. For help with data collection and related matters, I thank Dimka Atanassov, Jamie Fisher, Rachel Stults, Dorothy Ahn, and Robert Wilder. Thanks also are due to Dorothy Ahn for creating the visual stimuli used in Experiment 2.

References

Abbott, Barbara. 2008. Issues in the semantics and pragmatics of definite descriptions in english. In Jeannette K. Gundel & Nancy Hedberg (eds.) (eds.), *Reference: Interdisciplinary perspectives*, 61–72. Oxford: Oxford University Press.

Abrusán, Márta & Kriszta Szendrői. 2013. Experimenting with the king of France: Topics, verifiability, and definite descriptions. *Semantics & Pragmatics* 6(10). 1–43.

Altmann, Gerry & Mark Steedman. 1988. Interaction with context during human sentence processing. *Cognition* 30(3). 191–238.

Altmann, Gerry T.M. 1998. Ambiguity in sentence processing. *Trends in Cognitive Sciences* 2(4). 146–152.

Barr, Dale J., Roger Levy, Christoph Scheepers & Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68(3). 255–278. doi:10.1016/j.jml.2012.11.001.

Bates, Douglas M. 2005. Fitting linear mixed models in R. *R News* 5. 27–30.

Bott, Lewis & Ira A. Noveck. 2004. Some utterances are underinformative: The onset and time course of scalar inferences. *Journal of memory and language* 51(3). 437–457. doi:10.1016/j.jml.2004.05.006.

Burkhardt, Petra. 2006. Inferential bridging relations reveal distinct neural mechanisms: Evidence from event-related brain potentials. *Brain and Language* 98(2). 159–168.

Chemla, Emmanuel & Lewis Bott. 2013. Processing presuppositions: Dynamic semantics vs pragmatic enrichment. *Language and Cognitive Processes* 38(3). 241–260.

Clark, Herbert H. & William G. Chase. 1972. On the process of comparing sentences against pictures. *Cognitive Psychology* 3(3). 472–517.

Crain, Stephen & Mark Steedman. 1985. On not being led up the garden path. the use of context by the psychological parser. In David. R. Dowty, Lauri Karttunen & Arnold Zwicky (eds.), *Natural language parsing: Psychological, computational, and theoretical perspectives*, 320–358. Cambridge: Cambridge University Press.

Donnellan, Keith. 1966. Reference and definite descriptions. *Philosophical Review* 75(3). 281–304.

Elbourne, Paul D. 2013. *Definite descriptions*. Oxford: Oxford University Press.

Fiedler, Klaus, Eva Walther, Thomas Armbruster, Doris Fay & Uwe Naumann. 1996. Do you really know what you have seen? Intrusion errors and presuppositions effects on constructive memory. *Journal of Experimental Social Psychology* 32(5). 484–511.

von Fintel, Kai. 1994. *Restrictions on quantifier domains*: University of Massachusetts Amherst PhD thesis.

von Fintel, Kai. 2004. A minimal theory of adverbial quantification. In Hans Kamp & Barbara H. Partee (eds.), *Context-dependence in the analysis of linguistic meaning*, 137–175. Elsevier.

Frege, Gottlob. 1892. On sense and reference. In Peter Geach & Max Black (eds.), *Translations from the philosophical writings of gottlob frege*, 56–78. Oxford: Blackwell.

Garrod, Simon C. & Anthony J. Sanford. 1982. The mental representation of discourse in a focussed memory system: Implications for the interpretation of anaphoric noun phrases. *Journal of Semantics* 1(1). 21–41.

Gilbert, Daniel T. 1991. How mental systems believe. *American Psychologist* 46(2). 107–119.

Gilbert, Daniel T., Douglas S. Krull & Patrick S. Malone. 1990. Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology* 59(4). 601–613.

Grice, H.P. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Syntax and semantics*, vol. 3, 41–58. New York: Academic Press.

Grice, Paul. 1981. Presupositions and conversational implicature. In Peter Cole (ed.), *Radical pragmatics*, 183–198. New York: Academic Press.

Hawthorne, John & David Manley. 2012. *The reference book*. Oxford: Oxford University Press 1st edn.

Heim, Irene. 1982. *The semantics of definite and indefinite noun phrases*: University of Massachusetts Amherst PhD thesis.

Heim, Irene. 1983. On the projection problem for presuppositions. In M. Barlow, D. Flickinger & N. Wiegand (eds.), *Proceedings of WCCFL* 2, 114–125. Stanford University.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Malden and Oxford: Blackwell.

Heller, Daphna, Daniel Grodner & Michael K. Tanenhaus. 2008. The role of perspective in identifying domains of reference. *Cognition* 108(3). 831–836.

Horn, Laurence R. 1985. Metalinguistic negation and pragmatic ambiguity. *Language* 61(1). 121–174.

Klinedinst, Nathan. 2012. THCSP. Ms., UCL.

Kramer, Ruth & Kyle Rawlins. 2011. Polarity particles: An ellipsis account. In Suzi Lima, Kevin Mullin & Brian Smith (eds.) (eds.), *Proceedings of NELS 39*, vol. 2, 479–492. Amherst, MA: GLSA.

Kripke, Saul A. 1977. Speaker's reference and semantic reference. In P. French, T.E. Jr. Uehling & Howard K. Wettstein (eds.), *Contemporary perspectives in the philosophy of language*, 6–27. Minneapolis: University of Minnesota Press.

Lasersohn, Peter. 1993. Existence presuppositions and background knowledge. *Journal of Semantics* 10(2). 113–122.

Lewis, David. 1979. Scorekeeping in a language game. In Rainer Bäuerle, Urs Egli & Arnim von Stechow (eds.), *Semantics from different points of view*, 172–187. Berlin: Springer.

Loftus, Elizabeth F. 1975. Leading questions and the eyewitness report. *Cognitive Psychology* 7. 560–572.

Loftus, Elizabeth F., David G. Miller & Helen J. Burns. 1978. Semantic intergration of verbal information into a visual memory. *Journal of Experimental Psychology: Human Learning and Memory* 4(1). 19–31.

Loftus, Elizabeth F. & Guido Zanni. 1975. Eyewitness testimony: The influence of the wording of a question. *Bulletin of the Psychonomic Society* 5. 86–88.

Ludlow, Peter & Gabriel Segal. 2004. On a unitary semantical analysis for definite and indefinite descriptions. In Anne Bezuidenhout & Marga Reimer (eds.), *Descriptions and beyond*, 420–436. Oxford: Oxford University Press.

Murphy, Gregory. 1984. Establishing and accessing referents in discourse. *Memory & Cognition* 12(5). 489–497.

Reinhart, Tanya. 1981. Pragmatics and linguistics: An analysis of sentence topics. *Philosophica* 27. 53–93.

Roberts, Craige. 1996. Information structure in discourse: Towards an integrated formal theor of pragmatics. In J. H. Yoon & Andreas Kathol (eds.), *Papers in semantics* (OSU Working Papers in Linguistics 49), 91–136. OSU.

Roberts, Craige. 2003. Uniqueness in definite noun phrases. *Linguistics and Philosophy* 26(3). 287–350.

Romoli, Jacopo & Florian Schwarz. 2015. An experimental comparison between presuppositions and indirect scalar implicatures. In Florian Schwarz (ed.), *Experimental perspectives on presuppositions*, 215–240. Springer International Publishing.

Russell, Bertrand. 1905. On denoting. *Mind* 14. 479–493.

Schlenker, Philipe. 2008. Presupposition projection: The new debate. In Tova Friedman & Satoshi Ito (eds.), *Proceedings of SALT 18*, Ithaca, NY: CLC Publications.

Schlenker, Philippe. 2009. Local contexts. *Semantics and Pragmatics* 2. doi:10.3765/sp.2.3.

Schoubye, Anders J. 2010. Descriptions, truth value intuitions, and questions. *Linguistics and Philosophy* 32(6). 583–617.

Schwarz, Florian. 2009. *Two types of definites in natural language*. Amherst, MA: University of Massachusetts Amherst dissertation.

Schwarz, Florian. 2012. Domain restriction and discourse structure - evidence from online processing. Poster Presentation, 25th CUNY Sentence Processing Conference.

Schwarz, Florian. 2013. Maximality and definite plurals – experimental evidence. In Emanuel Chemla, Vincent Homer & Gregoire Winterstein (eds.), *Proceedings of Sinn und Bedeutung 17*, 509–526. http://semanticsarchive.net/sub2012/: Semantics Archive.

Schwarz, Florian & Sonja Tiemann. 2012. Presupposition processing - the case of German *wieder*. In Maria Aloni, Vadim Kimmelmann, Floris Roelofsen, Galit W. Sassoon, Katrin Schulz & Matthijs Westera (eds.), *Proceedings of the 18th amsterdam colloquium*, 200–209. Berlin: Springer.

Schwarz, Florian & Sonja Tiemann. 2013. The path of presupposition projection in processing – the case of conditionals. In Emanuel Chemla, Vincent Homer & Gregoire Winterstein (eds.), *Proceedings of Sinn und Bedeutung* 17, 509–526. Semantics Archive. http://semanticsarchive.net/sub2012/.

Sedivy, Julie & Greg N. Carlson. 2011. *Sold on language : how advertisers talk to you and what this says about you*. Hoboken: Wiley.

Sedivy, Julie C. 2003. Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of Psycholinguistic Research* 32(1). 3–23. doi:10.1023/a:1021928914454.

Simons, Mandy. 2001. On the conversational basis of some presuppositions. In Rachel Hastings, Brendan Jackson & Zsofia Zvolenszky (eds.), *Proceedings of SALT* 11, 431–448. Ithaca, NY: CLC Publications.

Singh, Raj, Evelina Fedorenko & Edward Gibson. 2013. Presupposition accommodation is costly only in implausible contexts. Manuscript, Carleton University & MIT.

Stalnaker, Robert. 1974. Pragmatic presuppositions. In Milton K. Munitz & Peter K. Unger (eds.), *Semantics and philosophy*, 197–213. New York: New York University Press.

Strawson, P. F. 1964. Intention and convention in speech acts. *The Philosophical Review* 73(4). 439–460.

Strawson, Peter F. 1950. On referring. *Mind* 59. 320–344.

Sudo, Yasutada. 2012. *On the semantics of phi features on pronouns*. Cambridge, MA: Massachusetts Institute of Technology PhD thesis.

Szabó, Zoltán G. 2000. Descriptions and uniqueness. *Philosophical Studies* 101(1). 29–57.

Szabó, Zoltán G. 2003. Definite descriptions without uniqueness: A reply to Abbott. *Philosophical Studies* 114(3). 279–291.