

# Overcoming the identity crisis: novel evidence for a referential theory of Verb Phrase Ellipsis\*

Till Poppels & Andrew Kehler  
UC San Diego

## 1 Introduction

It is well established that Verb Phrase Ellipsis (VPE), illustrated in (1), canonically exhibits a close correspondence between the ellipsis clause and the antecedent clause.

(1) The students had to study hard because the teacher didn't. [study hard]

This correspondence manifests itself in two ways. First, the acceptability of VPE appears to depend on the existence of some form of parallelism between the ellipsis clause and the antecedent clause. Second, when such parallelism constraints are met, the set of possible interpretations of the ellipsis clause tends to be highly restricted by the meaning of the antecedent clause. In the case of (1), for example, the meaning of the ellipsis clause must involve the meaning of the antecedent VP *study hard*; the ellipsis simply cannot be resolved to the proposition that the teachers didn't teach the material well, for instance, even though that would arguably result in a more plausible interpretation.

Broadly speaking, there are two ways in which these constraints have been captured in theories of VPE. The first has been to posit the existence of particular constraints on the felicitous use of VPE, ones that impose a requirement for identity between the antecedent and ellipsis clauses at some level of linguistic representation (most notably, at the level of syntax or semantics). We term such theories *IDENTITY ACCOUNTS*. The second is to treat VPE as a form of discourse reference, thereby assimilating its behavior to that of other familiar forms of reference, such as pronouns. We term such theories *REFERENTIAL ACCOUNTS*.

Associated with each type of theory is a set of well-known problems. Identity theories have to contend with cases in which VPE is acceptable and interpretable despite the lack of a suitable antecedent, which has motivated researchers to propose various rescue strategies to preserve identity. Referential theories, on the other hand, need to explain why various cases of VPE are infelicitous despite the fact that the required referent could plausibly be seen to be inferrable from the discourse context.

We begin the paper by surveying both theories and the predictions they make with respect to certain cases of VPE. We then present results from an acceptability

---

\*Thanks to Roger Levy as well as members of *Semantics Babble* and the *Experimental Syntax Lab* at UCSD for helpful comments and discussion. We gratefully acknowledge support from NSF grant BCS-1456081.

judgment experiment using a type of example not previously discussed in the literature, which, we argue, are only consistent with a referential account. Although the results ultimately lead us to advocate for the referential account, we conclude by discussing some of the remaining challenges for such theories.

## 2 Identity Theories of VPE

As noted above, the correspondence illustrated with (1) has inspired a class of identity theories of ellipsis, which posit that eliding linguistic material is only grammatical when it is identical (in some relevant way) to an antecedent VP. On one class of identity theory (Sag 1976, Merchant 2013a, *inter alia*), the ellipsis site contains linguistic material, and identity constraints license a speaker's choice to delete it or leave it unpronounced. On another class of identity theory, there is no unpronounced material at the ellipsis site (Dalrymple *et al.* 1991, Fiengo & May 1994, *inter alia*); identity constraints instead fall out from procedures that reconstruct a representation of the ellipsis clause from operations performed on a representation of the structure or meaning of the antecedent clause. Much of the previous literature on VPE has focused on arbitrating among different types of identity theories or determining at which level of representation identity is evaluated. Here we collapse different flavors of identity theories in order to focus on an important property they have in common, specifically that they reduce the constraints on VPE to the linguistic antecedent, without appealing to non-linguistic aspects of the context.

The attempt to reduce the context-dependency of VPE to an identity relation between the ellipsis site and the linguistic antecedent makes identity theories vulnerable to a variety of counterexamples that are acceptable and interpretable but lack the kind of antecedent VP deemed necessary. Consider (2), for instance.

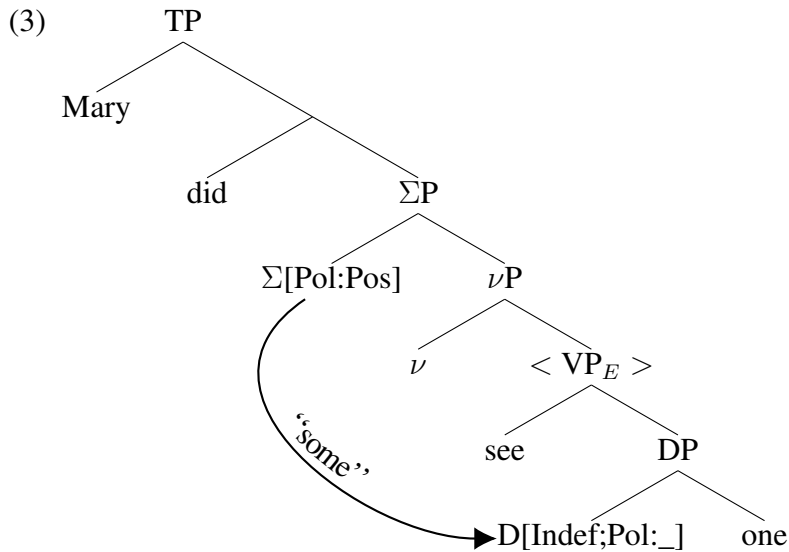
- (2) a. **Lexical mismatches:** John didn't see anyone, but Mary did ~~see someone~~. (Merchant 2013a)
- b. **Split antecedents:** John wanted to go to Bolivia and Mary wanted to go to Peru, but because it's expensive, neither of them can ~~go to Bolivia or Peru~~, respectively. (adapted from Webber 1978)
- c. **Exophora:** [Situational context: Two kids standing in front of a pool]  
First kid: I will ~~jump~~ if you will ~~jump~~. (Chao 1988, cited in Hardt 1993)

On a syntactic identity theory, the VP *see anyone* in example (2a) would be ungrammatical at the ellipsis site (*\*Mary did see anyone*). Instead, the VP required to derive the attested meaning would be *see someone*, but eliding it should be ungrammatical for lack of an identical antecedent. Similarly, in example (2b), there is no single VP that can serve as the antecedent, and the necessary VP is instead 'split' across two distinct clauses. Finally, example (2c) involves EXOPHORA, in which the meaning of the missing VP is recovered from the situational, and hence non-linguistic, context. Thus, all three examples in (2) showcase perfectly acceptable and interpretable VPE despite the fact that neither of them provide an antecedent VP that is identical to the elided material.

Various proposals have been offered to deal with cases in which VPE is acceptable despite the violation of identity, which can be roughly classified into two

categories: REPRESENTATIONAL approaches, which re-analyze mismatching linguistic elements in a way that preserves identity, and INFERRED-ANTECEDENT approaches, which invoke cognitive machinery to infer an identical antecedent if the linguistic context fails to provide one.

Representational approaches involve re-analyzing elements that appear non-identical across the antecedent and ellipsis sites. For example, syntactic identity can be preserved VP-internally if we analyze non-identical elements in terms of mismatching values that are assigned post-syntactically to otherwise identical abstract elements by agreement with a VP-external phrasal head. To illustrate, consider the following analysis of the lexical mismatch involving polarity items in (2a), which is due to Merchant (2013a):



The polarity item in the elided VP *see anyone* (as well as the one in the antecedent not shown here) is analyzed as an abstract lexical item at the level of syntax, which is lexicalized post-syntactically as either *some* (as in this example) or *any*, in agreement with the head of the VP-external polarity phrase. The apparent lexical mismatch between *any* in the elided VP and *some* in the antecedent is thus moved outside the elided constituent and attributed to the post-syntactic lexicalization process. This analysis thus establishes the isomorphism between antecedent and elided VP that syntactic identity theories require. Although we illustrate the logic behind the representational approach to non-identity with this specific example, similar strategies have been applied to a variety of other types of mismatches as well.<sup>1</sup>

Another strategy for reconciling identity theories with cases of non-identity involves postulating cognitive machinery that is capable of inferring a suitable antecedent when the context does not provide one. One example of this strategy is known as ANTECEDENT ACCOMMODATION (Fox 1999; Van Craenenbroeck 2012; Thoms 2015), which extends the formulation of identity to include a set of “accom-

<sup>1</sup>For instance, see Merchant (2013b) for an application of this analytic strategy to voice mismatches.

modatable” antecedents in addition to the one the context provides.<sup>2</sup> For example, Thoms (2015:ex. 51) posits the following mechanism for accommodating alternative antecedents, adapting similar machinery proposed by Katzir (2007) for generating alternative utterances to be used for deriving implicatures:

(4) **Accommodating alternative antecedents for ellipsis**

- a. A set of additional antecedents, Ad(A), may be accommodated on the basis of the original (overt) antecedent A.
- b. The members of Ad(A) are alternatives derived from A by
  - (i) deletion
  - (ii) contraction
  - (iii) substitution
- c. Complexity constraint: all members of Ad(A) must be at most as complex as the overt antecedent A.
- d. Semantic constraint: all members of Ad(A) must be semantically identical to the overt antecedent A.

Extending the identity requirement to the set of accommodatable antecedents Ad(A) makes it possible to capture certain mismatches, such as the polarity-item mismatch described above, so long as the mismatching elements can be replaced with ones that preserve identity. Since this is a powerful mechanism that, if left unrestricted, would be prone to overgeneration, Thoms (2015) stipulates additional constraints that require accommodated antecedents to be semantically identical to and at most as complex as the overtly provided antecedent. It is important to note at this point that this notion of accommodation is independent of discourse-level presupposition accommodation in the Lewisian sense (Lewis 1979). Antecedent accommodation operates on syntactic representations by deleting, contracting, or substituting words and phrases, whereas presupposition accommodation operates on discourse commitments and the common ground between interlocutors.

### **3 Overcoming the IDENTITY crisis: a referential theory of VPE**

On a referential account (Webber 1978, Hardt 1990), VPE is not associated with a set of ellipsis-specific rules that impose constraints on identity, but instead operates as a null pro-form. Hence, VPE interpretation is predicted to behave in a way similar to other pro-forms in the language. Crucial to contemporary theories of reference is the observation that reference cannot be aptly characterized by way of only a relationship among pieces of associated linguistic material, i.e. between a pro-form and its antecedent. Instead, the felicitous use of referring expressions depends on the status of the referent with respect to the hearer’s knowledge state, which includes (but is not limited to) their evolving mental model of the discourse. This mental model is referred to as the DISCOURSE MODEL, which contains a structured record of the entities and eventualities that have been introduced and the relationships that hold among them (Karttunen 1976, Webber 1983, inter alia). The felicity of a ref-

---

<sup>2</sup>Another influential proposal for inferring antecedents is the RECYCLING hypothesis (Arregui et al. 2006; Frazier 2013). While we focus on antecedent accommodation in this paper, all arguments extend straightforwardly to Recycling-based analyses.

erence might thus depend on factors such as whether the speaker believes that the hearer has prior knowledge of a referent, whether the referent has been mentioned previously in the discourse, whether the referent is situated in the immediate surroundings, how salient the speaker thinks the referent is in the hearer's mind, the level of topicality of the referent, and the degree of difficulty in accommodating a discourse model representation for a referent when one doesn't already exist.

On the referential theory, therefore, VPE succeeds by virtue of a suitable representation of the referent existing in, or being constructable from, the hearer's discourse model. Consider (1) again, repeated in (5), and augmented with the representation for the pro-form  $\phi$ :

- (5) The students had to study hard because the teacher didn't  $\phi$ . [study hard]

Among the discourse model representations generated as a result of the compositional semantic analysis of the sentence (which will also include representations for the meanings of noun phrases such as *the students* and *the teacher*) will be a representation for the meaning of the VP *study hard*. This discourse model representation thus provides a suitable referent for the subsequent VPE. The discourse model does not, however, contain the VP meaning necessary for deriving the more plausible proposition *the teacher didn't teach the students well*, which is therefore not a possible interpretation of the ellipsis clause.

Support for a referential theory of VPE comes from the fact that it patterns with other types of pronominal reference in a variety of ways. For instance, it predicts that, like pronouns, VPE should be able to refer cataphorically and obey the same constraints on co-reference that govern pronominal cataphora (Lakoff 1968; Jackendoff 1972). Consider (6), where subscripts are used to indicate co-reference:<sup>3</sup>

- (6) a. If McCain will  $\phi_i$ , Obama will [make a statement criticizing President Putin] <sub>$i$</sub> .  
b. If he <sub>$i$</sub>  makes a statement criticizing President Putin, Obama <sub>$i$</sub>  will make a fool of himself.  
c. \* Obama will  $\phi_i$ , if McCain will [make a statement criticizing President Putin] <sub>$i$</sub> .  
d. \* He <sub>$i$</sub>  will make a fool of himself, if Obama <sub>$i$</sub>  makes a statement criticizing President Putin.

Cataphoric VPE is permitted when it is embedded in a subordinate clause as in (6a), and the same is true of pronominal reference in (6b). Cataphora is not permitted, however, when the ellipsis is not embedded in a subordinate clause, which again holds for both VPE in (6c) and pronominal reference in (6d). On an ellipsis-specific identity analysis, this pattern can only be coincidental, but it is just what is predicted on a referential account.

Another property of pronominal reference is that while it has a strong preference for referents evoked from within the current or immediately preceding clause, it is also able to access ones evoked further back. VPE displays a similar pattern; Hardt (1990) reports that five percent of the examples in the Brown corpus have an

---

<sup>3</sup>Examples taken from Kehler (2017).

antecedent that is at least two sentences back in the discourse, as in example (7).

- (7) The thought came back, the one nagging at him these past four days. He tried to *stifle it*. But the words were forming. He knew he *couldn't*.

Likewise, the subject pronoun *he* in the final sentence in passage (7) locates its referent from two sentences back.

Another property of reference is that it allows for split antecedents. Specifically, pronouns can access referents that are evoked from multiple, distinct constituents in the discourse. Consider again (2b), adapted in (8).

- (8) John wanted to [go to Bolivia]<sub>i</sub> and Mary wanted to [go to Peru]<sub>i</sub>, but because it's expensive, neither of them can  $\phi_i$ . [go to Bolivia or Peru, respectively] (Webber 1978)

As discussed in Section 2, the meaning that  $\phi$  refers to depends simultaneously on two distinct antecedent VPs, *go to Bolivia* and *go to Peru*. Analogously, *them* successfully refers to Mary and John as a group, even though they, too, are introduced from distinct constituents. Even more tellingly, the pronoun *it* appears to depend the same two VPs as  $\phi$ : what is understood to be expensive is going to either Bolivia or Peru. This example thus demonstrates that both VPE and other discourse referential devices can refer to meanings that have been introduced by several conspiring, yet syntactically distinct, constituents.

Yet another property of reference is its ability to be exophoric, that is, to refer to an entity or eventuality that is situationally, and not linguistically, introduced, as in (9):

- (9) [Situational context: A course lecturer walks into class five minutes late]  
Student: There he is!

- (10) [Situational context: Two kids standing in front of a pool]  
First kid: I will  $\phi$  if you will  $\phi$ . [jump]

Similarly, in (2c), adapted above as (10), VPE retrieves its referent exophorically.<sup>4</sup>

Finally, and importantly, under certain conditions an anaphoric pronoun can access a referent that is not introduced directly by, but is nonetheless inferrable from, an antecedent expression. Consider (11).

- (11) Jean is a Frenchman, although he hasn't lived there for many years. (Ward et al. 1991)

Here, the pronoun *there* successfully refers to France, which has not been explicitly mentioned. However, the referent has a transparent semantic relationship to the meaning of the antecedent *Frenchman*, and hence is derivable from it. Similarly, Webber (1978) offers an example of a related sort involving VPE:

---

<sup>4</sup>Note that whereas pronouns and VPE can both refer exophorically in principle, exophoric VPE appears to be more constrained than exophora in the nominal domain (Schachter 1977; Hankamer 1978). The question remains as to whether this fact should cast doubt on the hypothesis that VPE is referential, or if there are independent discourse factors that can explain the difference. See Miller & Pullum (2013) for discussion and arguments for the second of these positions.

- (12) Martha and Irv wanted to dance with each other, but Martha couldn't, because her husband was here.

Whereas the antecedent clause does not contain a VP that denotes the intended referent – i.e., the meaning of *dance with Irv* – this meaning can be transparently inferred from the meaning of the antecedent clause.

The behaviors of VPE illustrated here are thus highly consistent with a referential account. Furthermore, some of these behaviors are among the biggest challenges for identity accounts of VPE, such as its ability to refer in split-antecedent or exophoric contexts. It is worth stressing, however, that these arguments have proceeded primarily by analogy: it is widely accepted that pronominal reference does not require an identity relation between a referring expression and its antecedent, and since VPE resembles pronominal reference in many respects, it stands to reason that VPE is not governed by identity constraints either. What we do not have is a predictive theory that can distinguish cases of successful reference from unsuccessful ones (particularly in cases of split or inferred antecedents, for example), for either pronominal reference or VPE. While we recognize developing such a predictive theory as an important long-term goal for discourse theorists, this is not the goal we pursue in this paper. Instead, we will aim to show that the referential approach is the one that is more likely to ultimately prove successful. To establish this, we sought to find a new type of example that is compatible with referential accounts of VPE, but is far out of reach of identity theories as well as the various rescue strategies posited to reconcile data that are problematic for such accounts.

For this purpose, we turn to the domain of indirect speech acts, which we illustrate first with an example involving pronominal reference. To set up the context, imagine that Susan and Tom are in their family room, night has just fallen after a warm day, and Susan is standing next to an open patio door at the time of the exchange in (13).

- (13) Tom: I'm cold.  
Susan: Sure, I'll close it in a second. [*it* = patio door]

Here, Susan has interpreted Tom's statement as an indirect speech act, specifically a request to close the patio door. Susan responds in accordance with this interpretation, and indeed considers the patio door to have been made salient and topical enough by Tom's utterance to be able to refer to it with the pronoun *it*, even though it has not been explicitly mentioned. Importantly, note that the reference may be perceived as slightly degraded – for example, compared to the use of *the patio door* instead of *it* – an intuition that is expected in light of the theory of reference, since a modicum of inferential work is required to accommodate a discourse model referent for the door. But it succeeds nonetheless.

Now consider a discourse of the sort shown in (14).

- (14) Spectator: Can I please see that card trick one more time?  
Magician: I'm sorry, I can't.

On perhaps the most natural interpretation given the context, the magician's utterance would be intended to mean *I'm sorry, I can't show you the card trick one*

*more time*, and not the one that identity would predict, i.e., *I'm sorry, I can't see the card trick one more time*. As we will argue below, the interpretation involving a card-trick *performing* event, rather than the card-trick *observing* event introduced by the antecedent VP, lies outside of the explanatory reach of the identity rescue strategies surveyed in Section 2. On the other hand, such a reading is compatible with a referential theory of VPE, in light of the fact that the spectator's utterance comes with an indirect speech act conveying a request that the magician perform the trick again. Hence, while not linguistically expressed, a performing event might be made accessible enough for subsequent reference, albeit perhaps with some level of degradedness, on analogy with the status of the patio door in example (13).

This sets the stage for the experiment described in the following section, whose purpose was two-fold: to test whether dialogues like (14) indeed make interpretations beyond the antecedent VP available, and to test whether such interpretations are relatively acceptable despite the extreme violation of identity that they entail. Towards this end, we couple canonical acceptability judgments with a novel experimental task, in which participants paraphrase the ellipsis site.

## 4 Methods

### 4.1 Participants and materials

We recruited 20 self-reported native English speakers from Amazon.com's Mechanical Turk, who were compensated according to the federal minimum wage of \$7.25/hour. They were presented with ten two-turn dialogues like the one in (14), repeated in (15) to illustrate the experimental manipulation. The first utterance contained an active-voice VP ( $V = \{see, get, hear, know, borrow\}$ ), and the second utterance employed auxiliary-focus VPE.

(15) **Spectator:** Can I please see that card trick one more time?

- a. **Magician:** I'm sorry, you can't. [NO-CHANGE variant]  
b. **Magician:** I'm sorry, I can't. [CHANGE variant]

The first utterance was designed to do two things: to provide a VP that could serve as the antecedent for subsequent ellipsis, and to communicate an indirect speech act by implicitly submitting a request—in this case, that a card trick be performed. As for the response, we manipulated in a within-item and within-participant design whether the pronominal subject of the ellipsis clause referred to the same person as the antecedent-clause subject (the NO-CHANGE condition) or to a different person (the CHANGE condition). Consequently, the antecedent VP was identical across conditions, but the antecedent-provided referent yielded a plausible interpretation only in the NO-CHANGE condition (*You can't see that card trick one more time* [you = the spectator]), and not in the CHANGE condition (*I can't see that card trick one more time* [I = the magician]).

The ten experimental items were interspersed with 20 non-elliptical and 20 elliptical filler dialogues, and the latter were used to establish upper- and lower-bound baselines for the acceptability and interpretation of run-of-the-mill VPE. Lower-bound elliptical fillers involved violations of various sorts, and lower-bound non-elliptical fillers involved extraction violations. Examples of filler items are given in



(16) below.<sup>5</sup>

(16) a. **Upper-bound elliptical filler:**

Can you see the remote control anywhere? —I can't, sorry.

b. **Lower-bound elliptical filler:**

Of course I believe in God.

—Even though a proof that God exists doesn't?

c. **Upper-bound non-elliptical filler:**

I can't hear what he's saying. —I don't care.

d. **Lower-bound non-elliptical filler:**

I'm telling you, I didn't get your email.

—That's impossible, whose did you get email?

## 4.2 Procedure

Dialogues were presented twice in separate blocks, each time in random order. In the first block, participants were asked to rate the second utterance in terms of its acceptability in context (cf. Fig. 1A). In the second block, they were asked to paraphrase the ellipsis site in a free production task (cf. Fig. 1B). The purpose of this task was to investigate what interpretations the ellipses received and to what extent those interpretations deviated from the meaning of the antecedent VP. To that end, the paraphrases were analyzed by extracting the main verb and comparing it to the head of the preceding VP to determine whether comprehenders' interpretation of the ellipsis site matched the preceding VP. Reducing this comparison to the verb itself greatly improved the tractability and transparency of the analysis of the free-production results from the paraphrase task.

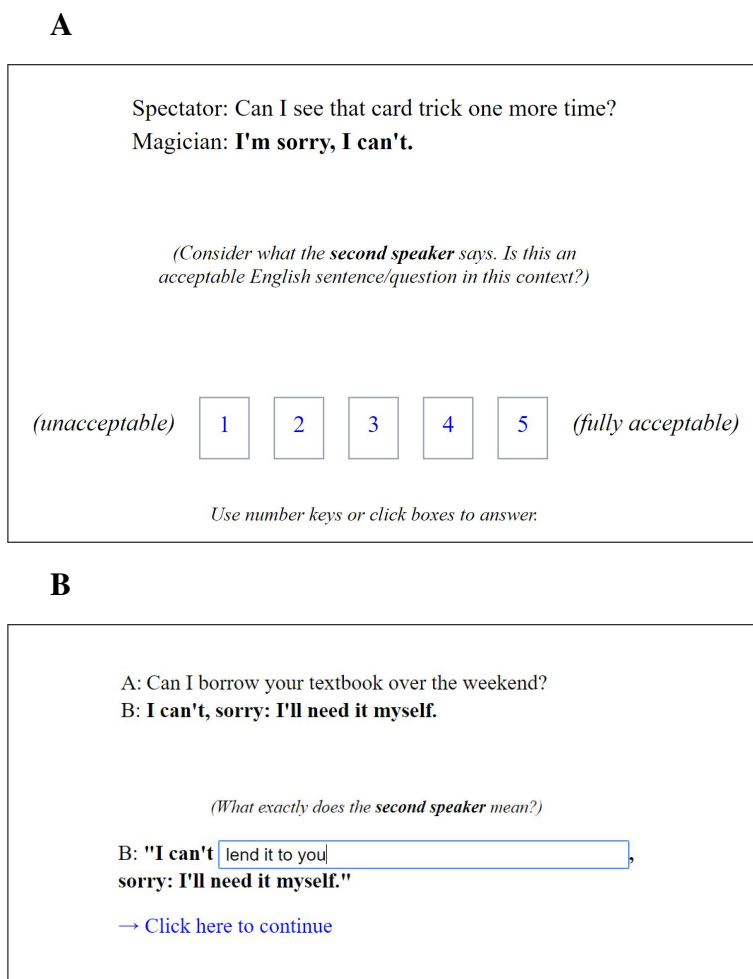
## 4.3 Analysis and predictions

If comprehenders interpret the dialogues in line with the grammatical constraints of identity theories, both NO-CHANGE and CHANGE items should be interpreted according to the antecedent-provided meaning. Consequently, we would expect paraphrases of the ellipsis site to use the antecedent verb to express that meaning to similar degrees in both conditions, and we would expect comparable levels of acceptability (modulo a possible reduction in acceptability of the CHANGE variant due to a reduction in plausibility) since items in both conditions observe identity. If, on the other hand, comprehenders adopt a different interpretation for CHANGE items (e.g. *I can't show you the card trick again* in (15b)), that should be reflected in a lower degree of overlap between antecedent and paraphrase verbs than in the NO-CHANGE condition. That sort of interpretation would entail an extreme mismatch between the antecedent VP and the ellipsis site, which should thus result in low acceptability, perhaps on a par with that of lower-bound elliptical fillers.

The referential theory of VPE makes different predictions. On that view, we expect the ellipsis site to be able to pick up inferred referents if the discourse context is sufficiently conducive. Thus, comprehenders should be free to adopt interpretations that deviate from the antecedent VP to the extent that our experimental items

---

<sup>5</sup>The full list of experimental and filler stimuli is available upon request; email [tillpoppels@gmail.com](mailto:tillpoppels@gmail.com).




**Figure 1:** Screen shots illustrating the (A) acceptability judgment task and the (B) paraphrase task. Both examples show CHANGE variants.

succeed in providing such a conducive context. Any difference in acceptability between CHANGE and NO-CHANGE items would be expected to be in proportion to the inferential work associated with accommodating the necessary referent. We further analyzed the variability associated with participants' verb choice in the paraphrase task because we suspected that, if comprehenders do infer VPE meanings beyond the one provided by the antecedent VP, those inferred meanings may be less precisely "carved out" than referents that were introduced by purely linguistic means. This prediction reflects the privileged status the referential theory assigns to antecedent-provided referents.

## 5 Results

The results are summarized in Figure 2. There was a strong tendency for CHANGE items to deviate in their interpretation from the preceding VP: only 5% of para-

phrases employed the antecedent verb, compared to 74% for their NO-CHANGE counterparts ( $\beta = 4.47, p < .001$ ; see Fig. 2B). (The rate of antecedent verb re-use in the NO-CHANGE condition was numerically, but not significantly, lower than that of upper-bound elliptical fillers, whose mean is shown by the dashed line.) Despite this deviation, CHANGE items were much more acceptable than lower-bound elliptical fillers ( $\beta = 1.39, p < 0.001$ ), and only slightly less acceptable than their NO-CHANGE counterparts ( $\beta = -0.22, p = 0.011$ ; cf. Fig. 2A).



**Figure 2:** Results from the acceptability judgment task (A) and paraphrase task (B and C). B shows the proportion of antecedent verb use; C shows variation (entropy) in verb choice. Dashed lines show the mean scores for elliptical fillers, exemplified in (16a) and (16b).

Finally, we analyzed the degree of uncertainty in participants’ verb choice during the paraphrase task, operationalized as entropy in the by-item distribution of verbs  $v \in V$ :

$$(17) \mathcal{H}(V) = - \sum_{v \in V} P(v) \log P(v)$$

CHANGE items were associated with significantly more uncertainty in paraphrase verb choice ( $t = 2.14, p = 0.048$ ), although that effect was relatively small (cf. Fig. 2C).

## 6 Discussion

Let us first consider the results from the perspective of identity theories. The results from the paraphrase task show that CHANGE items overwhelmingly received interpretations that deviated dramatically from the antecedent-provided meaning. Since they were nonetheless judged relatively acceptable, our results are problematic for identity theories in light of the extreme degree of mismatch between antecedent and elided VPs. For example, consider the results for the CHANGE variant of the card trick item discussed earlier. To derive the participants’ interpretations under identity, the elided material must involve either *show* or *do*, both of which featured in roughly 50% of paraphrases:

(18) **Spectator:** Can I please see that card trick one more time?

**Magician:** I'm sorry, I can't  $\left\{ \begin{array}{l} \text{show you the card trick again.} \quad (\approx 50\%) \\ \text{do the card trick again.} \quad (\approx 50\%) \end{array} \right.$

In both cases, the elided material deviates substantially from the antecedent VP, and together they exhibit two types of mismatch: both involve several lexical mismatches, including the verb itself, and the *show* variant additionally involves a mismatch in transitivity, with ditransitive argument structure in the ellipsis clause where the antecedent clause is transitive. To our knowledge this is the first report of VPE with this degree of mismatch that nonetheless achieves high levels of acceptability.

At this point we may ask whether these cases of non-identity can be accommodated under an identity analysis with the help of the rescue strategies surveyed in Section 2. Let us consider the representational approach first. Recall that this approach would involve reanalyzing all mismatching elements as covertly identical, which in this example will minimally require the following identities: *see* = *show*; transitive = ditransitive. Provided that such a reanalysis is possible at all (it is not clear to us how), this strategy will run into fatal problems once we consider a slightly broader range of data. For example, consider the following item along with the elided material that would be required to derive participants' interpretations under identity:

(19) **A:** Before Trump got elected, people demanded to see his tax returns, but he refused.

**B:** And now that he's president, I don't think

he ever will  $\left\{ \begin{array}{l} \text{release his tax returns.} \quad (\approx 63\%) \\ \text{show his tax returns.} \quad (\approx 27\%) \\ \text{provide his tax returns.} \quad (\approx 9\%) \end{array} \right.$

If we just focus on the verb, the representational approach to these cases of non-identity would now require not just *see* = *show* = *do*, but also *see* = *provide* and *see* = *release*, which by transitivity would result in unattested identities, such as *do* = *provide*. As a result, such an analysis would grossly overgenerate and, for example, falsely predict that *release the card trick* should be a possible interpretation of the CHANGE variant of the card trick item.

Inferred-antecedent strategies, such as Thoms' (2015) antecedent accommodation algorithm summarized in (4), face similar challenges when applied to our results. Under such an analysis, the elided material *may* differ from the overt antecedent if it is identical to some member of the set of accommodatable antecedents. Thus, if we allow the necessary lexical substitutions (and/or insertions and deletions) to take place, we may derive the attested interpretations under identity and explain the high degree of acceptability CHANGE items achieved. Note, however, that permitting the necessary lexical substitutions, such as substituting *show* for *see*, would violate the semantic constraint in (4d) that requires inferred antecedents to be semantically identical to the overt one. Further, other necessary operations, e.g. accommodating a ditransitive antecedent VP from the transitive overt one to derive the *show* interpretation of the card trick item, would further violate the complexity constraint in (4c). These constraints were specifically put in place to prevent overgeneration, and it is easy to see that permitting the operations that would be necessary to derive the attested interpretations would additionally generate a wide

range of other interpretations that are unattested.

Thus, the cases of VPE we examined in our experiment represent a class counterexamples to identity theories of VPE that neither the representational nor the inferred-antecedent rescue strategy can defuse convincingly. By contrast, the finding that VPE is amenable to accommodated referents in conducive contexts is consistent with a referential account. Particularly the finding that the inferences in question are sensitive to considerations of pragmatic inference and world knowledge is unsurprising given that this is a well-known property of discourse accommodation in other referential domains.

We may ask, however, why CHANGE items were reliably degraded compared to their NO-CHANGE counterparts (Fig. 2A). The referential theory we sketched in Section 3 predicts that whenever a reference requires accommodating inferences, its felicity depends on the inferential burden comprehenders face in resolving it. Recall in this regard the example provided in (13), in which the intended referent of the pronoun *it* is made available by an indirect speech act requesting that the patio door be closed. While *it* refers successfully, it may be degraded compared to other referring expressions, such as the definite NP *the patio door*, which would require less inferential work to accommodate the discourse conditions it presupposes. Likewise, while *there* is straightforwardly resolved to *France* in (11), repeated in (20a) below, the analogous inference appears somewhat less straightforward in (20b).

- (20) a. Jean is a Frenchman, although he hasn't lived there for many years.  
b. Tim is a Dutchman, although he hasn't lived there for many years.

This demonstrates that the felicity of referring expressions that require accommodation depends on the difficulty of performing those inferences in the context. Since the difference between CHANGE and NO-CHANGE items in our experiment was that the former did require such accommodation while the latter did not, it is not surprising that those inferences incurred some cost in acceptability.

## 7 Conclusion

We have introduced a type of VPE example not previously discussed in the literature in which the meaning of the ellipsis clause is contributed not by the antecedent clause directly, but rather an indirect speech act that it conveys. We have presented experimental evidence that such interpretations do indeed arise when they are supported by the discourse context, and that they achieve relatively high levels of acceptability. This finding is inconsistent with identity theories of VPE as well as proposed rescue strategies for cases that involve mismatch. Furthermore, our examples exhibit the symptoms of discourse accommodation that are familiar from the literature on reference, and hence point us in the direction of a referential theory of VPE.

These results may seem surprising in light of the fact that, as discussed at the outset of the paper, VPE appears to be immune to plausibility-based considerations in many other cases. For example, recall (5), repeated below as (21):

- (21) The students had to study hard because the teacher didn't  $\phi$ . [study hard]

Despite the implausibility of the antecedent-provided interpretation and the apparent availability of a more plausible alternative reading, that reading is not a possible interpretation of the ellipsis clause. Whereas this result follows trivially from identity accounts of VPE, a referential theory has to provide an account for why the necessary inference to accommodate the reference is not supported in this context.

The answer, we believe, bears on a salient disanalogy between cases like (21) and the passages used in our experiment such as (15b). In (21), there are many possible interpretations for the VPE that result in a plausible outcome – *the teacher didn't teach them well, the teacher didn't provide a sample exam, the teacher didn't assure them that the test would be relatively straightforward*, etc. As with any anaphoric expression, context must provide a way of distinguishing a particular interpretation among other possible ones. This is why a sentence like *John said that he studied hard*, spoken discourse initially and outside of a special context (e.g., one in which another male is situationally present), can only be interpreted with *he* referring to John: the addressee is not free to resolve the pronoun to some other male referent, even if plausibility factors point away from John as the referent (e.g., John is notorious for never studying). Similarly, the context in (21) only distinguishes one possible referent from other possible ones – i.e., the one that was linguistically evoked – and hence VPE only succeeds when this is the intended referent.

The situation is different for (15b), however. In this case, the referent required for the non-identity reading is inferred, but it is not inferred in service of satisfying the referential requirements of VPE. It is instead inferred as a by-product of recognizing the spectator's question as an indirect request to show (or do, or perform) the trick. This has the effect of introducing this potential interpretation into the discourse model as well, with the result that there are now two possible interpretations that are distinguished from all of the others: one introduced by linguistic evocation (*see the trick*), and one resulting from the inference of an indirect speech act (*show the trick*). As our results demonstrate, both are indeed available for subsequent reference with VPE, albeit with the second, inferred interpretation being slightly less accessible.

This observation may explain another type of case that may at first seem puzzling for the referential account, illustrated in (22):

(22) a. After the test I wasn't sure if I had passed or not.

As it turns out,	{	I did.	[pass][*fail]
		I didn't.	[pass][*fail]
		it's less likely than I thought.	[pass][*fail]

Even though the context makes both passing and failing a salient possibility, only the meaning of the antecedent VP—in this case *passing* the exam—is a possible referent for subsequent VPE. While this observation may be interpreted as *prima facie* evidence for identity accounts of VPE, this cannot be the right story, as the same constraint exists for propositional anaphora: *it* in *it's less likely than I thought* is also limited to the meaning introduced by the antecedent verb. As in the cases surveyed above, in such a context the addressee is faced with the problem of individuating an interpretation from all possible ones, and here one interpretation has

been linguistically mentioned, i.e. *passed*. As we would expect, therefore, the judgments reverse if we replace *failed* with *passed* in (22).

To conclude, we have argued that the data examined in this paper provide substantial evidence for a referential theory of VPE as opposed to identity accounts. At the same time, they also highlight the importance of further investigating the contextual conditions under which referent can be accessed via VPE. We believe this to be an important goal for future research focused on building increasingly predictive theories of discourse reference in general and of VPE in particular.

## References

- Arregui, A., C. Clifton, L. Frazier, & K. Moulton. 2006. Processing elided verb phrases with flawed antecedents: The recycling hypothesis. Journal of memory and language 55.232–246.
- Chao, W. 1988. On ellipsis. Dissertations-G.
- Dalrymple, M., S. M. Shieber, & F. Pereira. 1991. Ellipsis and higher-order unification. Linguistics and Philosophy 14.399–452.
- Fiengo, R., & R. May. 1994. Indices and Identity. Cambridge, MA: MIT Press. Linguistic Inquiry Monograph Twenty-Four.
- Fox, D. 1999. Focus, parallelism and accommodation. In Semantics and Linguistic Theory, volume 9, 70–90.
- Frazier, L. 2013. A recycling approach to processing ellipsis. Cheng and Corver (2013) 485–501.
- Hankamer, J. 1978. On the nontransformational derivation of some null VP anaphors. Linguistic Inquiry 9.66–74.
- Hardt, D. 1990. A corpus-based survey of VP ellipsis. University of Pennsylvania.
- Hardt, D. Verb phrase ellipsis: Form, meaning, and processing. University of Pennsylvania dissertation.
- Jackendoff, R. S. 1972. Semantic Interpretation in Generative Grammar. MIT Press.
- Karttunen, L. 1976. Discourse referents. In Syntax and Semantics VII: Notes from the Linguistic Underground, ed. by J. McCawley, 363–386. New York: Academic Press.
- Katzir, R. 2007. Structurally-defined alternatives. Linguistics and Philosophy 30.669–690.
- Kehler, A. 2017. Ellipsis and discourse. In Handbook of Ellipsis, ed. by J. van Craenenbroeck & T. Temmerman. Oxford University Press. To appear.
- Lakoff, G. 1968. Pronouns and reference. In Syntax and Semantics, Volume 7: Notes from the Linguistic Underground, ed. by J. McCawley, 275–335 (1976). New York: Academic Press. Originally distributed by Indiana University Linguistics Club, 1968.
- Lewis, D. 1979. Scorekeeping in a language game. Journal of Philosophical Logic 8.339–359.
- Merchant, J. 2013a. Polarity items under ellipsis. Diagnosing syntax 441–462.
- Merchant, J. 2013b. Voice and ellipsis. Linguistic Inquiry 44.77–108.
- Miller, P., & G. K. Pullum. 2013. Exophoric VP ellipsis. The core and the periphery: Data-driven perspectives on syntax inspired by Ivan A. Sag 5–32.
- Sag, I. Deletion and logical form. Massachusetts Institute of Technology dissertation.
- Schachter, P. 1977. Does she or doesn't she? Linguistic Inquiry 8.762–767.
- Thoms, G. 2015. Syntactic identity, parallelism and accommodated antecedents. Lingua 166.172–198.
- Van Craenenbroeck, J. 2012. Ellipsis, identity, and accommodation. Ms. Hogeschool-Universiteit Brussel.
- Ward, G., R. Sproat, & G. McKoon. 1991. A pragmatic analysis of so-called anaphoric islands. Language 67.439–474.
- Webber, B. L. 1978. A formal approach to discourse anaphora. Technical report, Bolt Beranek And Newman, Inc, Cambridge, MA.
- Webber, B. L. 1983. So what can we talk about now? In Computational Models of Discourse, ed. by M. Brady & R. C. Berwick, 331–371. Cambridge, MA: The MIT Press.