Palaz, Bilge & Bruening, Benjamin & Tollan, Rebecca. 2025. Sluicing and subject islands: An experimental approach. *Glossa: a journal of general linguistics* 10(1). pp. 1–30. DOI: https://doi.org/10.16995/glossa.11054

Open Library of Humanities

Sluicing and subject islands: An experimental approach

Bilge Palaz, University of Delaware, US, bilgeplz@udel.edu

Benjamin Bruening, University of Delaware, US, bruening@udel.edu

Rebecca Tollan, University of Delaware, US, rtollan@udel.edu

This paper investigates how sluicing in English interacts with islandhood (Ross 1967), focusing on how the subject island constraint applies under both regular sluicing (with an indefinite correlate) and contrast sluicing (with a focused correlate). We conducted three experiments. Experiment 1 examined what the structure of the elided clause is, by having participants choose a possible continuation. We found different patterns for regular sluicing versus contrast sluicing, with contrast sluicing exhibiting a strong preference for syntactic parallelism with the antecedent clause. Regular sluicing, in contrast, preferred a non-parallel continuation (a copular clause), consistent with evasion approaches to island repair under sluicing (e.g., Barros et al. 2014). Next, Experiment 2 examined the sensitivity of sluicing to island effects. Strikingly, contrast sluicing did not demonstrate notable acceptability degradations when subject islands were involved (contra predictions made by Merchant 2008; Barros et al. 2014; Griffiths & Lipták 2014), but we did find an overall degradation in contrast sluicing in comparison to regular sluicing. Because this outcome challenges previous assertions regarding the sensitivity of contrast sluicing to island constraints, Experiment 3 asked whether the subject island effect is truly at play in embedded *wh* questions. We found that – like with matrix questions but unlike with relative clauses (Abeillé et al. 2020) – subject island effects *do* emerge in embedded interrogatives. Since evasion is not a possible approach for contrast sluicing, whatever the source of the subject island effect is, it disappears when the relevant structure is not pronounced.

1 Introduction

Sluicing, first described by Ross (1969), is a type of clausal ellipsis where a *wh*-phrase by itself is understood to correspond to a full interrogative clause. Sluicing is attested both in matrix clauses, as in B's response in (1), and in embedded clauses, as in (2). The focus of this paper will be on constructions like that in (2), which constitute 72% of all natural instances of sluicing (Anand, Hardt & McCloskey 2021).

(1) A: John was talking to someone on the phone.
B: Really? Who?

(2) John was talking to *someone* on the phone, but I don't know *who*.

Conventionally, the first clause (i.e., "John was talking to someone on the phone") is called the *antecedent clause*, and the indefinite (i.e., "someone") in the antecedent clause is referred as the *correlate*. The *wh*-phrase in the second clause (i.e., "who") is called the *wh-remnant*. The correlate and the remnant in (1) and (2) correspond to the same person in the discourse whose identity is unknown. The wh-remnant is interpreted as a full clause, as *Who was John talking to on the phone?* in (1), and as *who John was talking to on the phone* in (2).

There are different kinds of sluicing that are identified based on the properties of the correlate. We will refer to sluicing constructions where the correlate is an indefinite, as in (1) and (2), as *regular sluicing.* The indefinite correlate in regular sluicing can be a pronominal element (1–2), or it can be a full Noun Phrase (NP) as in (3).

(3) John was talking to *a student* on the phone, but I don't know *who/which student*.

The correlate can also be a definite expression (Chung et al. 1995; Merchant 2001; 2008). Sluicing constructions where a definite correlate bears prosodic prominence (e.g., it is contrastively focused) are referred to as *contrast sluicing*.

(4) a. Abby speaks *Greek*, but I don't know *what other language(s)*.
b. *Beth* was there, but you'll never guess *who else*.
(Merchant 2001: 3 (1c))

In contrast sluicing, the *wh*-remnant involves an extra linguistic element such as "other" or "else", commonly referred to as an *exceptive modifier* (see Barros 2014b). The exceptive modifier ensures that the referent(s) of the remnant are the other candidate(s) in discourse that is/are distinct from the correlate.

Contrast sluicing is claimed not to be uniformly acceptable in every kind of linguistic environment (e.g., Merchant 2008; Barros et al. 2014; Griffiths & Lipták 2014). Contrast sluicing is observed to be unacceptable, for instance, when the correlate ("the cake" in 5a) resides in an

if-clause (bracketed), as shown in (5a). This is unlike regular sluicing with a correlate in the same structural configuration in (5b) (e.g., Ross 1969; Merchant 2001).

(5) a. *Ben will be angry [if you don't try *the cake*], but I don't know *what else*.
(Barros et al. 2014: 20 (90))
b. Ben will be angry [if you don't try *an appetizer*], but I don't know *which appetizer*.

Interestingly, these types of environments are the same environments as those which disallow *wh* dependency formation in constructions such as *wh* questions. Environments that block *wh* dependencies are known as *islands*. Islands have long been taken to be structural domains out of which *displacement* or *movement* of a phrase is not allowed (e.g., Chomsky 1973; Ross 1967). Displacement causes the sentence to be unacceptable, as in (6).

(6) *W*hich appetizer*$_1$ will Ben be angry [if you don't try t_1]?

In comparison with (6), the acceptability pattern in (5a) versus (5b) appears to be puzzling. The unacceptability of (5a) makes sense, if the elided clause is isomorphic to the antecedent clause, since the wh-remnant must have undergone extraction out of an island (…*what else Ben will be angry [if you don't try]*). But then we would expect (5b) to be unacceptable, as well. Griffiths & Lipták (2014) generalize the contrast between (5a) and (5b), and state that contrastive ellipsis like contrast sluicing is island sensitive, but non-contrastive ellipsis like regular sluicing is insensitive to islands. However, recent experimental work has challenged this, for instance, by finding contrastive *stripping* to be island insensitive (Yoshida et al. 2019), which we detail in Section 4. No previous work has experimentally investigated the acceptability pattern in (5a–b)—contrast sluicing versus regular sluicing—, a gap which we seek to fill in the present study.

Both regular sluicing and contrast sluicing have generated a lot of interest in the linguistic literature, because – unlike (6) – there are no pronounced strings of words after the *wh*-remnant in these constructions. In other words, meaning is somehow computed from silence. To account for how this interpretative process plays out, there are two strands of approaches: *structural* and *non-structural* approaches. Broadly speaking, non-structural approaches consider sluicing to involve sentential "fragments" with no covert syntax (e.g., Culicover & Jackendoff 1995; Ginzburg & Sag 2000; Barker 2013; Nykiel & Kim 2022). The interpretation of sluicing is achieved via featural properties of the syntactic node that dominates the *wh*-word (Culicover & Jackendoff 1995; Ginzburg & Sag 2000), or via an anaphoric relation between the antecedent clause and a null element in the ellipsis site (Chung et al. 1995; Barker 2013).

The other approach, the structural approach, maintains that there *is* covert syntax at the ellipsis site (e.g., Ross 1969; Lasnik 2001; Merchant 2001). For instance, there might be a structure matching the antecedent clause, as was illustrated in example (6) above. One popular way of modeling this, pioneered by Merchant (2001), is to give the head C^0 an ellipsis feature. This feature requires non-pronunciation of C^0's complement (i.e., TP, as shown in gray ink in (7)).

(7) John was talking to someone, but I don't [$_{VP}$ know [$_{CP}$ who$_1$ [$_{TP}$ John was talking to t_1]]].

In this approach, sluicing deletes a clausal constituent and leaves behind only the *wh*-phrase that escaped from the ellipsis site via movement (Ross 1969; Merchant 2001; and much other work).

As shown in (5a–b), regular and contrast sluicing seem to respond differently to islands: regular sluicing appears to be insensitive to islands, as in (5b), but contrast sluicing is sensitive to islands (as in 5a). To explain the insensitivity of regular sluicing to islands, some structural approaches have argued for an *ellipsis repair* mechanism where the full isomorphic structure is present but unpronounced, and this lack of pronunciation repairs the island violation (Ross 1969; Merchant 2001; 2008). Others have proposed an *evasion* strategy (Barros et al. 2014), where only a simpler structure, called a "short source," is present, and movement in this simple elided "short source" structure does not involve movement out of an island. Importantly, both repair and evasion approaches predict contrast sluicing to be island sensitive but differ in their explanation for why (we lay out the full details of these two approaches in Section 2).

With this backdrop in mind, the primary goal of this paper is to evaluate the predictions of the non-structural, repair, and evasion approaches using novel experimental data. We do so by examining how regular sluicing and contrast sluicing behave when their antecedent is embedded either inside a subject NP or an object NP, like with "a monster" in (8).

(8) a. **Correlate inside subject:**
[Stories about a monster] $_{SUBJECT}$ terrified John, but I don't know which monster.
b. **Correlate inside object:**
John heard [stories about a monster] $_{OBJECT}$, but I don't know which monster.

Extraction out of object NPs as in (9b) is generally considered to be grammatical, but extraction out of subject NPs like "stories about" in (9a) is not. (9a) constitutes a type of island known as a SUBJECT ISLAND (Ross 1967).

(9) a. **Extraction from a subject**: *Which monster$_1$ did [stories about t_1] terrify John?
b. **Extraction from an object**: Which monster$_1$ did John hear [stories about t_1]?

Subject islands have been a topic of widespread debate in linguistics and cognitive psychology, independently of sluicing. This debate centers around the question of whether subject island violations arise due to formal constraints in the syntax proper (e.g., Chomsky 1973; Huang 1982; Stepanov 2007), or instead, whether the syntax itself allows the extraction, but the resulting output is judged to be ill-formed because it creates a meaning which is anomalous in most kinds of discourse contexts (e.g., Erteschik-Shir 1973; Goldberg 2006; Chaves & Dery 2019; Abeillé et al. 2020). In light of this ongoing debate, a secondary goal of our current study is to lay out what kinds of data from sluicing the two competing approaches must take into account in order to be successful.

The paper proceeds in the following way. First, Section 2 presents more background on how the repair and evasion approaches explain island (in)sensitivity in clausal ellipsis (with a focus on subject islands). Section 3 presents Experiment 1, which examines what the nature of the unpronounced structure in sluicing is. Our results suggest that the ellipsis source differs in regular and contrast sluicing: contrast sluicing does not allow a short source, while regular sluicing does. Considering this strict syntactic isomorphism in contrast sluicing but not in regular sluicing, we next test subject island effects in these constructions. Experiment 2, presented in Section 4, quantitatively evaluates how regular and contrast sluicing respond to subject islands in English. Contrary to predictions (e.g., Merchant 2008), we find no island effect in contrast sluicing, just like in regular sluicing. In view of this finding (i.e., insensitivity of contrast sluicing to islands), we further examine subject islands in Experiment 3 as presented in Section 5. Experiment 3 explores subject island effects in embedded wh-questions, which are taken to be the pronounced counterparts of sluicing. The findings verify that the subject island effect is active in embedded wh-questions. Since contrast sluicing does not have a short source available to it (Experiment 1), this means that ellipsis must be repairing the island effect. The implications of these findings will be discussed further in Section 6.

2 Background: Islandhood under sluicing and how it has been accounted for

If one assumes that sluicing is a form of regular *wh*-movement, differing only in that the rest of the clause after the wh-phrase is unpronounced, then sluicing is expected to be constrained by the same set of rules that regulate *wh*-movement. Hence, an island effect is expected to be found in sluicing when the antecedent clause hosts a correlate inside an island (i.e., the sentence should be unacceptable). Starting with Ross (1969), sluicing has been extensively explored in relation to various kinds of islands, like the complex NP constraint and the adjunct island constraint as exemplified in (10) (island domains are bracketed).

(10) a. **Complex NP (relative clause) island:**
*Which Balkan language$_1$ does Abby want to hire the candidate [who speaks t_1]?
b. **Adjunct island:**
*Which appetizer$_1$ will Ben be angry [if you don't try t_1]?

Interestingly, although the structural approach expects island effects in sluicing, they have in fact been observed to be absent (e.g., Ross 1969; Merchant 2001). An example involving a complex NP object island is shown with a regular sluicing sentence in (11). In the structural approach, the wh-remnant is taken to have moved out of the exact same domain as in (10a) above (shown in gray in (11)). However, while (10a) is unacceptable, (11) is perfectly acceptable.

(11) *Regular sluicing (island insensitive)*
Abby wants to hire a candidate who speaks a Balkan language, but I don't know **which Balkan language**$_1$ [Abby wants to hire a candidate [who speaks t_1]].

Interestingly, this insensitivity to islands is claimed not to hold of contrast sluicing (e.g., Merchant 2008; Barros et al. 2014; Griffiths & Lipták 2014). Contrast sluicing is judged to be unacceptable when the correlate (e.g., "Greek" in (12)) is embedded in a complex NP object.

(12) *Contrast sluicing (island sensitive)*
*Abby wants to hire a candidate who speaks Greek, but I don't know **which other language**$_1$ [Abby wants to hire a candidate [who speaks t_1]].
(adapted from Merchant (2008): 148 (52a))

Two accounts have been proposed so far to explain the observed pattern of island (in)sensitivity under sluicing (11 vs 12). These accounts differ in their view of syntactic parallelism (i.e., structural matching) under sluicing. The *repair approach* (Chomsky 1972; Lasnik 2001; Merchant 2001; 2008) argues that the elided clause involves a structure isomorphic to the antecedent clause in both regular and contrast sluicing. This means the elided clause involves an island when the antecedent clause has one, as was shown in gray in (11) and (12). According to the repair approach, island violations occur due to PF (Phonological Form) constraints (e.g., Merchant 2001). The ellipsis operation is proposed to repair island violations via non-pronunciation (i.e., deletion at PF). Repair succeeds and sluicing becomes insensitive to the island effect in the case of regular sluicing (11). When the ellipsis is of the contrastive type, such as contrast sluicing (12), however, the sentence is unacceptable due to constraints on movement: the correlate is not allowed to move to the highest clausal node at LF and therefore fails to license the ellipsis (Merchant 2008); or, the correlate is banned from being extracted from an island at Logical Form (LF) (Griffiths & Lipták 2014).

The *evasion approach* adopts a more flexible view of the structural matching that is required to hold between the elided clause and the antecedent clause (Vicente 2008; Barros 2014a; 2014b; Barros et al. 2014; Griffiths 2019; Wu 2022). Syntactic mismatch is permitted because ellipsis can take place when a semantico-pragmatic identity condition is met according to this strand of approaches. Importantly for our purposes, these identity conditions are loose enough to allow structural non-parallelism between the antecedent clause and the elided clause. Accordingly, the elided clause can fail to match the syntax of the antecedent clause in regular sluicing; for instance, it can contain a copular clause, as shown in (13).

(13) Abby wants to hire a candidate who speaks another language, but I don't know **which language**$_1$ [that is t_1]].

The evasion approach, hence, naturally explains the acceptability of regular sluicing in relation to islands. There is no island violation in (13): note that pronouncing the material in gray in (13) is acceptable, in contrast with pronouncing the gray material in (11).

As for contrast sluicing, the evasion approach argues that the elided clause has to be syntactically isomorphic to its antecedent clause in contrast sluicing, otherwise, infelicity results due to independent focus-background matching constraints (Barros 2014a; Weir 2014; Griffiths 2019). This required matching forces the elided clause to contain an island if the antecedent clause has one, which then leads to unacceptability when the wh-remnant is extracted from it, as in (14).

(14) *Abby wants to hire a candidate who speaks Greek, but I don't know **which other language**$_1$ [Abby wants to hire a candidate [who speaks t_1]]. [= (12)]

Non-structural approaches to ellipsis predict no island effects whatsoever (Culicover & Jackendoff 1995; Ginzburg & Sag 2000; Barker 2013; Nykiel & Kim 2022), because there is simply no elided structure present at all – instead, the intended meaning of the ellipsis is reconstructed through featural properties of the syntactic node that dominates the *wh*-word (Culicover & Jackendoff 1995; Ginzburg & Sag 2000), or via an anaphoric relation between the antecedent clause and a null element (i.e., *e*) in the ellipsis site (Chung et al. 1995; Barker 2013). These approaches face a challenge from informal judgements like those in (12) and (14), where it appears that island effects *do* emerge in contrast sluicing. Thus, it is important to establish precisely what the empirical landscape of islandhood under regular and contrast sluicing looks like through fully controlled acceptability judgements.

To set the scene for investigating subject islands, we first provide new insights into what the ellipsis source is in regular and contrast sluicing by taking one step back from islandhood. In our first experiment (Section 3), we use a sentence continuation study to examine what kind of material is elided from the position after the wh-remnant in regular and contrast sluicing. We then use the results from this experiment to specify what predictions each approach (non-structural, repair, and evasion) makes regarding the effect subject islands should have on regular sluicing and contrast sluicing.

3 Experiment 1: Investigating the source under sluicing

The purpose of this experiment is to investigate what the structure of the unpronounced material may be in regular sluicing and in contrast sluicing. Formal research takes the (felicity of) potential non-elliptical continuations into consideration when deciding what the elided material must be (Merchant 2001; Barros 2014a; 2014b; Barros et al. 2014; among others). No prior study

has experimentally investigated such heuristic approaches that rely on native speaker intuitions about "what the sluice would sound like had it not been elided" (Barros 2014b: 4). According to Barros (2014b: 4):

> … the heuristic method of checking speaker intuitions about plausible continuations has proven to be useful in many domains in the investigation of ellipsis. It is, in fact, this intuition that motivates the isomorphism hypothesis to begin with, as simple examples of sluicing … are intuitively synonymous with a non-elliptical paraphrase that is the Wh-question version of the antecedent. Speaker intuitions about the meaning of the sluice, in particular, are an unquestionably valid source of empirical evidence in constraining the hypothesis space for theories of the E-site's content.

In the same spirit, we adopt the view that a pronounced felicitous continuation will mirror the structure that undergoes deletion under sluicing. Hence, we ask experimental participants to pick among possible continuations in order to tap into intuitions about what the elided material is likely to be.

We conducted a three-alternative forced-choice sentence completion experiment where participants had to read a sentence and choose the best continuation from three options. In a 2 × 2 (within-subjects, within items) design, we crossed two factors. The first was the sluice type: the stimuli involved either regular sluicing sentences (e.g., *a(n) x* in the first conjunct, *which x* in the second conjunct), or contrast sluicing sentences (e.g., *the x* in the first conjunct, *which other x* in the second conjunct). The second manipulation was the grammatical position of the NP correlate, as subject or direct object (NP_{subj}; NP_{obj}). The correlate was <u>not</u> embedded in any larger structure (like an island). This allowed us to assess the availability of each continuation independently of the presence of an island. The grammatical role manipulation was included in case there was an effect which might be relevant in testing subject islands in Experiment 2. This 2 × 2 paradigm is shown in **Table 1** (top part).

After each sentence in the experiment, three options were presented as a potential continuation as shown in **Table 1**, bottom part (note that these options do not have island violations, because the correlate NP is not embedded in any larger NP structure). Of the three continuation options, one involved a copular clause in the form of a pronoun "that", followed by a copular verb (i.e., "that is", "that was"). In evasion approaches, copular clauses are considered to be one of the alternative non-isomorphic structures that can appear in the ellipsis site under regular sluicing. They are thought not to be available in contrast sluicing, because of the independent focus-background constraint on contrastive focus (see section 2). Indeed, the two native speaker authors of this study judged the copular continuation to be unacceptable in

the contrast sluicing condition, in line with what has been reported in the literature (Merchant 2001; Vincente 2008; Barros 2014a; 2014b; Barros et al. 2014; Griffiths 2019; a.o.). Our goal was to verify this experimentally. In the present study, "that" was used as the subject in copular clauses rather than "it" to prevent ambiguity; "it" in this environment could also be the subject of a reduced cleft (Barros et al. 2014) rather than a simple copular clause. Using "it" rather than "that" could have introduced undesirable complexity to the stimuli; hence, the non-isomorphic copular continuation option consistently involved "that". We highlight that "that" is interpreted as an e-type pronoun in this environment (Merchant 2001; Barros et al. 2014): its interpretation is anaphoric to the meaning of the antecedent clause, which basically corresponds to the meaning of a relative clause headed by a definite NP derived from the correlate (e.g., *the superhero that amazed Mike*).

The second continuation option involved a syntactically isomorphic clause that matched the structure in the antecedent clause. Because antecedent clauses were created as simple English sentences with a transitive verb and its core arguments, the isomorphic continuation options had the same transitive verb as in the antecedent clause in addition to that verb's internal or external NP argument, both repeated as they appeared in the antecedent clause. The third option was a grammatically ill-formed continuation (which served as an attention-check). This experimental paradigm is shown in **Table 1**.

Condition	**Stimulus**	
Regular-Subject	A superhero amazed Mike, but his mother cannot recall which superhero…	
Contrast-Subject	Captain America amazed Mike, but his mother cannot recall which other superhero…	
Regular-Object	Mike loved a superhero, but his mother cannot recall which superhero…	
Contrast-Object	Mike loved Captain America, but his mother cannot recall which other superhero…	
	Continuation options	a) …that was. *Copular (non-isomorphic) continuation*
		b) …amazed Mike (b'. ...Mike loved.) *Isomorphic continuation*
		c) …the fun were. *Ungrammatical distractor*

Table 1*:* Sample stimuli in four conditions crossing sluice type (contrast vs regular sluicing) and the position of the correlate (subject vs object), followed by the three options presented.

3.1 Method

3.1.1 Materials

We created 16 item sets like those in **Table 1**, rotated in a Latin-square (each participant saw each item in only one of the 4 conditions, and saw a total of 4 items in each condition). These stimuli were coupled with eight filler sentences. The fillers included complex English sentences conjoined by "but" and ending with an adjunct wh-word, "why" or "when," as in (15). Similar to the critical stimuli, three potential continuation options presented for filler sentences involved a copular continuation as well as a non-copular continuation and an ungrammatical distractor.

(15) **A Sample Filler**
A school trip is coming up, but parents still don't know when…
a) …that is. *Copular continuation.*
b) …the students are leaving. *Non-copular continuation.*
c) …arriving the bus. *Ungrammatical distractor.*

The resulting 24 items were presented to participants in a pseudorandomized order: no more than three critical items were presented in a row, and no adjacent critical items were of the same condition. The presentation order of the three options was rotated across the experiment to counterbalance possible biases associated with order of mention.

The predictions for each theoretical stance on sluicing are as follows: First, repair approaches predict that the isomorphic continuation will be chosen as a response in both regular and contrast sluicing (i.e., no difference between the conditions). Evasion approaches predict that the copula response may be chosen in regular sluicing, but not in contrast sluicing (i.e., evasion approaches predict a difference between the two types of sluicing). As for non-structural approaches, which posit *no* structure at the ellipsis site, it is not entirely clear what they would predict. It must be possible to go from ellipsis to a pronounced correspondent in these approaches (when asked to clarify, e.g.), but it is not clear whether that pronounced correspondent would be something isomorphic to a linguistic antecedent for the ellipsis, or might be something else, like a copular clause. Without any considerations favoring one over the other, we might expect roughly 50–50 (i.e., at chance) selection of isomorphic or copula continuation (but no difference between regular and contrast sluicing). Finally, none of the approaches predict any difference between subjects and objects, as far as we can tell.

3.1.2 Participants

Forty native speakers of English were recruited on Amazon's Mechanical Turk (MTurk) to complete the study online. Participants were "master" workers who were independently qualified by MTurk based on their high-quality performance on previous tasks. The location was

restricted to the United States. Participants self-reported as L1 speakers of English. Compensation was not contingent on language background; all participants were paid 2 dollars in exchange for their participation.

3.1.3 Procedure

The experiment was created using the online software PCIbex Farm (Zehr & Schwarz 2018) and conducted on MTurk. Participants read and accepted a disclaimer approved by the University of Delaware's ethics board before starting the experiment. Participants were instructed to read sentences that would appear on their screen and pick the option that best completes each sentence. Three practice trials were included at the beginning of the experiment. No feedback was given. The task took approximately 10 minutes to complete.

3.2 Statistical analysis

Trials where the participants chose the ungrammatical distractor option (less than 1% of all trials) were removed from analysis. This left us with 637 critical trials for analysis. For the purposes of analysis, our dependent variable was whether participants picked the *isomorphic* continuation (coded as "1") or not (coded as "0"). We analyzed the output using two complementary approaches. First, we fitted a 2×2 logistic mixed-effects regression model with crossed random effects for participants and items for analysis (Baayen et al. 2008). The analysis was conducted in *R* version 4.3.2 using the *lme4* package (Bates et al. 2015; R Core Team 2021). The predictors *position* and *sluice type* were sum coded. We report the output for the model with the maximal random effect structure justified by the design that would allow for model convergence (Barr et al. 2013). All p-values were calculated by Satterthwaite approximation, using the *lmerTest* package (Kuznetsova et al. 2017).

Second, we complemented this frequentist analysis by computing Bayes Factors (BFs) to quantify the degree of support in favor of rejecting or accepting the null hypothesis. We implemented this step because both the repair approach and the non-structural approach do not predict a difference between regular and contrast sluicing conditions – in other words, the outcome should be a null (i.e., non-significant) main effect which, in a frequentist statistical paradigm, cannot be interpreted either as evidence in favor of the null hypothesis or against an alternative hypothesis. BFs were computed using the *brms* package in R (Bürkner 2017) by comparing the marginal likelihood of a model with the relevant effect against marginal likelihood of a model without it. For each model, we ran four Monte Carlo Markov chains, with 10,000 samples in each (the first 4,000 were discarded as part of warm-up), with the same random effect structure that we used for the mixed-effects regression model. We used weakly informative priors, following Avetisyan et al. (2020): we used a prior of N (0, 10) for

the intercept and priors of N (0, 1) for all predictors and interactions. The full 2 × 2 model is reported in supplementary files; we report only the BFs in the text here along with the frequentist outputs. BF_{10}'s > 3.0 (or < .33) demonstrate moderate to substantial evidence in favor of (or against) the alternative hypothesis over the null hypothesis. BF_{10}'s > 10.0 (or < .01) demonstrate strong to very strong evidence (Jeffreys 1998).

3.3 Results

The results are shown in **Figure 1**. Note that, since the dependent variable of the present study is binary, the remaining proportion of choices – the reverse of the pattern in **Figure 1** – corresponds to the copular continuation responses.

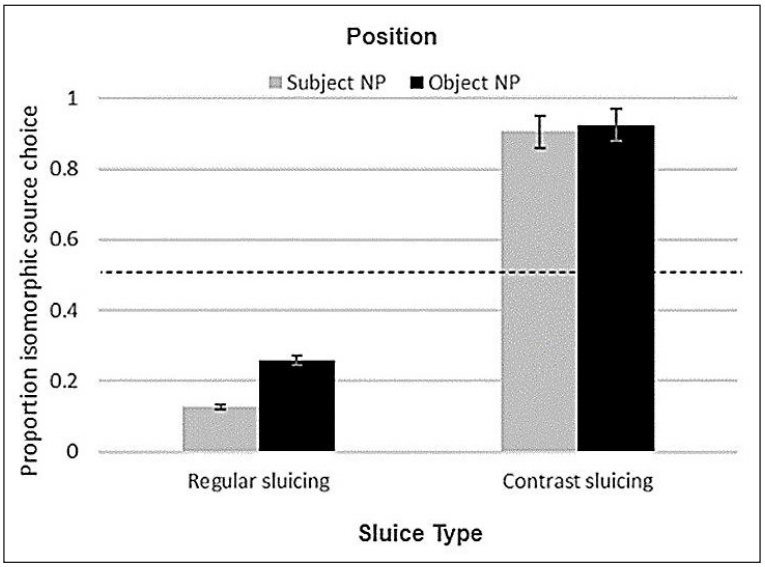

Figure 1: Proportion of choice of isomorphic continuation across contrast sluicing and regular sluicing. The dashed line represents chance level (i.e., 50%).

Sluice type was a significant predictor of participants' source choice: the isomorphic continuation was chosen significantly more in contrast sluicing conditions compared with regular sluicing (0.92 vs 0.2; $\beta = -6.34$, $SE = 0.87$, $z = -7.29$, $p < .0001$; $BF_{10} = 9 \times 10^7$). There was also a significant effect of position in that the isomorphic continuation was chosen significantly more with object correlates compared with subject correlates (0.58 vs 0.53; $\beta = -0.83$, $SE = 0.32$, $z = -2.56$, $p = .01$; $BF_{10} = 27$). The interaction of sluice type and position was marginal ($\beta = -1.17$, $SE = 0.65$, $z = -1.81$, $p = .07$), but the Bayes Factor analysis revealed that we have at best only weak evidence to believe that this marginal interaction corresponds to a meaningful effect ($BF_{10} = 1.82$). Thus, we leave it aside for present purposes.

Lastly, to evaluate whether the proportion of choice of isomorphic source was significantly different from chance level (i.e., 50%, see the dashed line in **Figure 1**) for each of the two sluicing levels, we conducted a further analysis by recoding the predictors (following West et al. 1996) and fitted a separate model for each level, focusing on the model intercept: because 0 corresponds to 0.5 (or 50%) in logit space, a significant positive intercept means that the response is statistically above chance, and a significant negative intercept means that it is below chance. Here, we found that the isomorphic continuation option was picked *above* chance level in contrast sluicing conditions (0.92; $\beta = 4.18$, $SE = 0.80$, $z = 5.17$, $p < .001$), but *below* chance level in regular sluicing. The pattern was in the opposite direction as indicated by the negative values of the intercept of the model (0.2; $\beta = –2.18$, $SE = 0.42$, $z = –5.2$, $p < .001$). In sum, a short (copula) continuation was preferred for sentences with regular sluicing, and an isomorphic continuation was preferred for sentences with contrast sluicing. We also found a higher rate of isomorphic continuation choice when the antecedent was in object position compared with subject position.

3.4 Discussion

We found that syntactically parallel (i.e., isomorphic) continuations were preferred in sentences with contrast sluicing (chosen 92% of the time). In contrast, the preferred continuation with regular sluicing was a copular construction (chosen 80% of the time). This provides experimental evidence for the common assertion in the literature that the elided structure in contrast sluicing must be isomorphic to the structure of its antecedent clause. Given our findings, we conclude that it is correct that a short source is unavailable in contrast sluicing. Contrast sluicing requires that the elided clause be syntactically isomorphic to the antecedent clause. In contrast, we found a preference for the copular continuation in regular sluicing. This is consistent with the evasion approach, which proposes that the elided clause in regular sluicing can be non-isomorphic with its antecedent. Note also that this preference is not just a general preference for the copular continuation, since we found a different preference on the filler items. In filler trials, involving sentences like those in (15), participants chose the non-copula continuation (e.g., 15b) 72% of the time (i.e., above chance level; $\beta = 2$, $SE = 1$, $z = 2$, $p = .044$).

We see two possible interpretations of our findings. First, we could assume that participants' preferences regarding pronounced continuations reflect the mental representation that they construct in sluicing. On this assumption, our findings are strong evidence for the evasion approach (Barros 2014a; Weir 2014; Griffiths 2019) and against the repair approach (Chomsky 1972; Lasnik 2001; Merchant 2001; 2008). Our experimental findings indicate that a short copular source is available (and preferred) for regular sluicing but is not available for contrast sluicing.

Second, we could reject the assumption that participants' preferences regarding pronounced continuations reflect the mental representation of an elided clause. They could reflect preferences holding solely over pronounced material. On this interpretation, participants might have a general preference to avoid repetition in pronunciation. This makes them overwhelmingly choose the copular continuation when it is available (with regular sluicing). This continuation must not be available with contrast sluicing, since the isomorphic continuation that was so strongly favored with contrast sluicing does involve complete repetition. But note that on this interpretation, we cannot conclude anything about the structure of the elided clause in regular sluicing, because the observed preference in the results is only about pronounced material, on this interpretation. We can, however, conclude that the copular short source is not available for contrast sluicing.

On either interpretation, then, we conclude that contrast sluicing requires complete syntactic isomorphism. We cannot definitively conclude anything about regular sluicing, although we believe our results are most compatible with the evasion approach. It is possible that participants' preferences regarding pronounced continuations do not reflect their mental representations of sluicing. If they do not, then our results are compatible with the repair approach, and potentially with non-structural approaches (Culicover & Jackendoff 1995; Ginzburg & Sag 2000; Barker 2013; Nykiel & Kim 2022). These might also have predicted at-chance responses for both regular and contrast sluicing, which is not what we found.

To summarize, our findings indicate that contrast sluicing requires complete syntactic isomorphism. Contrast sluicing does not permit a short (copular) source. Regular sluicing does permit a short (copular) continuation; this suggests but does not definitively show that the elided clause in regular sluicing can be non-isomorphic to its antecedent.

Our second finding was that there was a significant contrast between sentences with subject and object correlates. Having an object correlate increased the likelihood of choosing an isomorphic continuation compared with a subject correlate. Stated another way, the copula source was chosen more when the correlate was a subject than when it was an object. This contrast was not predicted by any of the approaches to clausal ellipsis. We speculate that the difference might arise from the desire to avoid complete repetition. In the subject condition, the isomorphic continuation completely repeats the antecedent clause (*A superhero amazed Mike, but his mother cannot recall which superhero amazed Mike*). In contrast, in the object condition, object extraction changes the word order and avoids complete repetition (*Mike loved a superhero, but his mother cannot recall which superhero Mike loved*). The desire to avoid repetition may therefore have been stronger in the subject condition, leading participants to choose the copular continuation more frequently. There are other possible explanations for the contrast, as well; we will have to leave exploring this unexpected finding to future research.

Equipped with these results, we now turn to subject islands in regular and contrast sluicing. In line with the evasion approach (which our results are most compatible with), we predict an island effect in contrast sluicing but not in regular sluicing. This is also what the informal judgements in the literature have indicated.

4 Experiment 2: Subject island (in)sensitivity in sluicing

Recall that regular sluicing and contrast sluicing have been reported to interact with islands differently. Regular sluicing with an indefinite correlate is characterized as being insensitive to islands; however, the island effect is argued to be active in contrast sluicing where the correlate and the remnant are contrasted (Merchant 2008). This pattern has been claimed to be a property of clausal ellipsis in general: fragment answers and stripping have also been claimed to be sensitive to islands (Merchant 2005; 2008; Griffiths & Lipták 2014). The common factor for island sensitivity is argued to be "contrastiveness" (see Griffiths & Lipták 2014). However, some contrastive fragment answers involving islands have been informally reported to be acceptable, as shown in (16) (see Barros 2014b; Weir 2014; contra Griffiths & Lipták 2014; this example involves a relative clause island).

(16) A: Do they grant scholarships to students that study *Spanish*?
B: No, *French*.

Moreover, some experimental studies have found mixed results regarding the island sensitivity of *stripping*, a type of clausal ellipsis which is commonly analyzed as akin to sluicing and fragment answers (Potter 2017; Yoshida et al. 2019; Wu 2022). Potter (2017) suggests that contrastive stripping, as illustrated in (17a), is "partially" sensitive to island effects (e.g., rated 4.8/7, in comparison to its baseline (17b), which is rated 5.2/7). But on the other hand, Yoshida et al. (2019) found that the acceptability of contrast sluicing sentences like (17a) was almost as high as the baseline in (17b) (i.e., 5.02 vs 5.21). They conclude that contrast stripping is, in fact, *not* sensitive to islands.

(17) *Contrastive Stripping ((a) involves a Relative Clause Island, (b) is control)*
a. Joe: Jordan mocked the salesman [$_{RC}$ who sold a car *to an elderly customer*].
Bill: No, *to Ashley*.
b. Joe: Jordan implied that the salesman sold a car *to an elderly customer*.
Bill: No, *to Ashley*.

(Potter 2017: 97 – 98 (120(ii)))

The empirical evidence for island effects under clausal ellipsis is controversial. The current paper aims to experimentally investigate how regular and contrast sluicing interact with subject islands. We chose to investigate subject islands for two reasons: First, no prior study has

quantified the behavior of subject islands under sluicing; and second, with subject islands, it is possible to isolate the effect of the island by comparing extraction from subject with extraction from object, both of which involve very similar structural complexity but only one of which involves an island. We conducted an acceptability rating study with a 3 × 2 factorial design, crossing Sluice Type with NP position. In the NP position manipulation, the correlate is embedded either inside a subject ("NP_{subj}"), from which extraction should give rise to an island effect, or inside an object ("NP_{obj}"), which serves as a non-island baseline. Sluice Type had 3 levels. An example item set is shown in **Table 2.** The Contrast condition consists of sluicing sentences that contrast a definite correlate with a remnant that includes an exceptive modifier ("which other x"). The Regular condition includes regular sluicing sentences with an indefinite correlate and a *wh*-remnant ("which x"). We included a baseline level for Sluice type, which had an embedded clause with no wh movement in it (called "NoSluice"). The NoSluice condition involves sentences conjoined by the subordinator *but* and followed by an articulated embedded clause (no ellipsis).

	NP_{subj} (potential subject island)	**NP_{obj} (no island, baseline)**
NoSluice	[The costume of a superhero] used to amaze Megan's children, but she cannot recall the popular villain.	Megan's children used to love [the costume of a superhero], but she cannot recall the popular villain.
Regular	[The costume of a superhero] used to amaze Megan's children, but she cannot recall which superhero.	Megan's children used to love [the costume of a superhero], but she cannot recall which superhero.
Contrast	[The costume of Captain America] used to amaze Megan's children, but she cannot recall which other superhero.	Megan's children used to love [the costume of Captain America], but she cannot recall which other superhero.

Table 2: Sample stimuli for Experiment 2.

4.1 Methods

4.1.1 Materials

24 critical item sets were created and distributed across six lists according to a Latin-Square. Correlates always followed the preposition "of" (in order to obviate any orthogonal effects of using different prepositions across items). 36 filler sentences were added to the critical stimuli. Half of the fillers were complex sentences conjoined by "but", and the other half were of the same fashion plus they involved an embedded adjunct *wh*-word. None of the fillers involved sluicing or any other form of ellipsis. In total, 18 fillers were grammatical sentences, whereas the

remaining 18 were ungrammatical sentences (ungrammaticality was triggered due to argument structure or word order violations). The presentation order of the stimuli was pseudorandomized and filler sentences were interspersed. No more than two critical items were presented adjacent to one another, and neither of two adjacent critical items were of the same condition.

4.1.2 Participants

72 participants completed this experiment. 36 were recruited on Amazon's Mechanical Turk (MTurk), and the other 36 were undergraduate students at the University of Delaware who completed the task via a weblink provided to them by the experimenter; they were compensated with 1 course credit in Linguistics. Participants recruited through MTurk were "master" workers, with an IP address from within the United States. At the end of the survey, they were asked about their language background (compensation was not contingent on their response). They were compensated with $2 US.

4.1.3 Procedure

The experiment was created and conducted on the online software PC Ibex Farm (Zehr & Schwarz 2018). Participants read and accepted a consent form approved by the University of Delaware Human Subjects Review Board before starting the experiment. In the experiment, participants were instructed to read complex English sentences that would appear on a computer screen and rate their acceptability on a seven-point scale (1-unacceptable, 7-acceptable). The experiment took approximately 10 minutes to complete.

4.2 Results

We report the results of seventy self-reported L1 English speakers (1680 critical trials); two participants were excluded because they did not report English as their L1. The mean of the ratings of the grammatical filler sentences was 6.10 (SD = 1.30); the mean of the ratings of the ungrammatical fillers was 1.55 (SD = 0.98). This indicates that the acceptability scale was used reasonably.

Sentences in the Contrast condition received lower ratings (4.16) compared with the Regular (5.54) and No Sluice (5.31) conditions, as shown in **Figure 2**. Overall, sentences in NP_{subj} conditions, where the correlate is embedded in an NP in subject position, were rated lower than the ones in NP_{obj} conditions (4.86 vs 5.15). Among the three levels of Sluice Type, the difference between NP_{subj} and NP_{obj} was largest for contrast sluicing (3.91 vs 4.40).

The raw ratings were normalized via z-score transformation prior to statistical analysis (with means and standard deviations estimated by participant). The z-score transformed data was fitted to a linear-mixed effects regression model with crossed random effects for participants and items (Baayen et al. 2008), implemented in *lme4* in R (Bates et al. 2015). The factor of Position

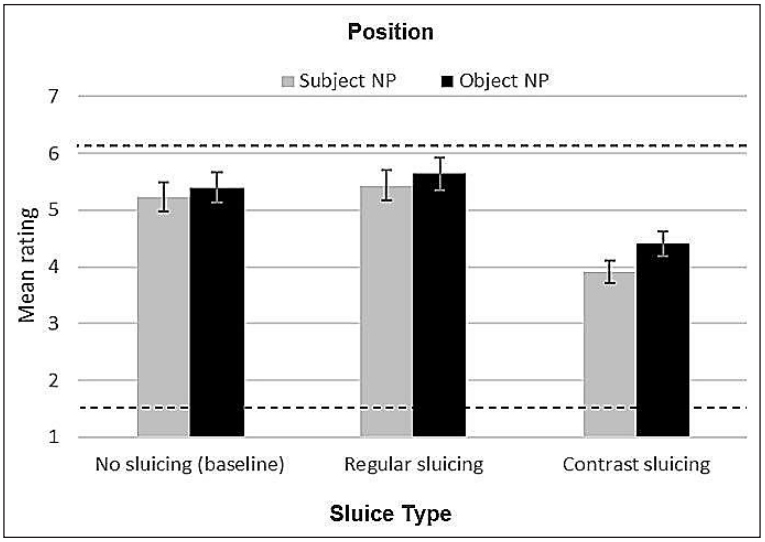

Figure 2: Mean ratings across No Sluicing, Regular Sluicing, and Contrast Sluicing conditions for Subject NPs (gray bars) and for Object NPs (black bars). Error bars represent 95% confidence intervals. The dashed line at the top represents the mean ratings to the grammatical fillers (6.10). The dashed line below is for the ungrammatical fillers (1.55).

was sum coded. The three levels of Sluice Type were coded with centered Helmert contrasts: The first contrast compared the No Sluice condition (coefficient: 2/3) with the two sluicing conditions (regular and contrast), pooled (coefficient: –1/3 for each). The second contrast compared Regular Sluicing (coefficient: –1/2) with Contrast Sluicing (coefficient: +1/2; the coefficient for No Sluice in this comparison was 0). We followed the same statistical protocols as for Experiment 1, likewise supplementing our frequentist results with Bayes Factors.

The results of the linear mixed-effects regression model are shown in **Table 3**. Our analysis revealed a significant effect of Position, which manifested as lower ratings to sentences when the correlate was embedded in an NP in subject position as compared to object position (4.86

EFFECT	*β*	*SE*	*t*	*p*
Position	**–0.13**	**0.04**	**–3.36**	**.003**
NoSluice vs. (Regular + Contrast)	**0.2**	**0.07**	**2.90**	**.007**
Regular vs. Contrast	**–0.61**	**0.05**	**–12.01**	**< .0001**
NoSluice vs. (Regular + Contrast) × Position	0.076	0.07	1.1	.27
Regular vs. Contrast × Position	–0.12	0.08	–1.6	.12

Table 3: Results of the linear-mixed effects regression model, Experiment 2. Significant effects are bolded.

vs 5.15; $\beta = -0.13$, $SE = 0.04$, $t = -3.36$, $p = .003$; $BF_{10} = 4.11$). No Sluice conditions were also rated as significantly better than the two sluicing conditions, pooled (5.54 vs 4.74, $\beta = 0.2$, $SE = 0.07$, $t = 2.9$, $p = .007$; $BF_{10} = 3$). Furthermore, Contrast Sluicing was overall rated significantly lower than Regular Sluicing (4.16 vs 5.31; $\beta = -0.61$, $SE = 0.05$, $t = -12.01$, $p < .0001$; $BF_{10} = 2.2 \times 10^{11}$).

Interestingly, however, neither interaction was significant ($ps > .12$), meaning we find no evidence that the effect of Position (subject vs object) changed according to Sluice Type. In fact, the Bayes Factor analyses indicated that we have moderate to strong evidence for accepting the null hypothesis; namely, that the effect of position is the same for non-sluicing as for regular and contrast sluicing pooled ($BF_{10} = .09$) and the same for regular sluicing and for contrast sluicing ($BF_{10} = .25$). Thus, we find that there is no subject island effect in either regular sluicing or, surprisingly, in contrast sluicing (although contrast sluicing is rated worse than regular sluicing generally).

4.3 Discussion

This study is the first to quantitatively examine subject island effects under sluicing. The findings only partially support the judgements reported in the literature: As reported, there is no subject island effect in regular sluicing. However – contrary to claims in the syntactic literature (e.g., Merchant 2008; Barros et al. 2014; Griffiths & Lipták 2014) – there is also no subject island effect in contrast sluicing. Contrast sluicing sentences were rated significantly lower than regular sluicing sentences overall, but this was not specific to subject NP islands. Our finding that contrast sluicing is not sensitive to island effects is however consistent with some informal judgments in the literature (e.g., Barros 2014b; Weir 2014), which report contrastive fragments to be acceptable with islands (see (16)), and the findings of Yoshida et al. (2019) in contrastive stripping settings (see (17)). Taken together, these findings challenge the generalization that contrastive ellipsis is island sensitive. To reconcile these findings with the informally reported judgments of sluicing, we suggest that the previous judgements of island sensitivity of contrast sluicing may reflect a more general degradation in contrast sluicing compared with regular sluicing, which we did find in our experiment. It seems likely that this general degradation is amplified with any kind of additional structural complexity, like that involved in islands, but we do not think it is specific to islands.[1]

Let us now consider how the three approaches to island sensitivity fare in predicting our results. As described above, the evasion approach assumes that evasion is not available

[1] Recall that the mean rating to contrast sluicing sentences was 4.16, and the mean rating to ungrammatical filler sentences was 1.55 in this experiment. It is therefore unlikely that the contrast sluicing sentences are subject to a floor effect, which would potentially mask a potential island effect if it were present. We thank an anonymous reviewer for raising this possibility.

to contrast sluicing, and so contrast sluicing should show an effect of islands while regular sluicing should not. This is not what we found, even though the results from Experiment 1 were most consistent with the evasion approach. Experiment 1 confirmed that contrast sluicing requires a completely isomorphic elided clause, which leads to the expectation of an island effect when the antecedent clause includes an island. However, this is not what we found in Experiment 2. Experiment 2 found no evidence of an island effect, only an overall degradation in contrast sluicing. The repair approach proposes that island repair is possible in regular sluicing but not in contrast sluicing, and so it too incorrectly predicts an effect of islands in contrast sluicing but not in regular sluicing. Overall, our findings best support the third type of approach, a non-structural approach to ellipsis (Culicover & Jackendoff 1995; Ginzburg & Sag 2000; Barker 2013; Nykiel & Kim 2022), because this is the only approach that, as currently formulated, predicts no island effect across-the-board. As we discussed in the previous section, however, pursuing a non-structural approach is problematic in light of our Experiment 1: There, we found evidence that an isomorphic structure *is* generally required for contrast sluicing sentences. In sum, the results of our Experiments 1 and 2 combined are not directly predicted by any of the three main approaches (non-structural, repair, evasion), as they are currently framed. One could amend the repair approach to say that island repair is possible for contrast sluicing as well as regular sluicing. It would be harder to amend the evasion approach to claim that evasion is possible with contrast sluicing, because our Experiment 1 found that it is not.

Alternatively, one might reason that contrast sluicing *is* island sensitive but that our Experiment 2 failed to find island sensitivity because we looked specifically at *subject* islands. Notably, the underlying source of the subject island effect has recently been called into question, with work by Abeillé et al. (2020) having established that extraction out of a subject only gives rise to an island effect in *wh* questions, but not in relative clauses. Abeillé et al. argue that this difference across construction types favors a discourse-based approach to subject islands, because critically, *wh* questions and relative clauses differ in their discourse functions but exhibit the same type of extraction (i.e., A-bar movement to spec, CP). Therefore, one possibility is that the elided material in sluicing constructions behaves like relative clauses, in that no island effect is present even when the material – an embedded *wh* interrogative construction – is overtly pronounced. Importantly, Abeillé et al. (2020) did not look at embedded wh questions – they compared matrix wh questions with relative clauses. Thus, our next goal is to establish how embedded wh questions pattern with respect to subject islands: like relative clauses (showing no island effect) or like matrix *wh* questions (showing an island effect)? Testing this is the aim of our Experiment 3, detailed in the next section.

5 Experiment 3: Subject island effects in embedded *wh* questions

The goal of this experiment is to determine how the elided material in sluicing patterns with respect to the subject island effect, when it is pronounced in the form of an embedded *wh* question as in the bolded part of the example in (18).

(18) [The costume of a superhero] used to amaze the children, but **Megan cannot recall which superhero [the costume of _] used to amaze the children.**

Recent work by Abeillé et al. (2020) has shown that subject island violations are not judged uniformly. In matrix wh-questions, subject island violations are judged to be unacceptable. However, in relative clauses, subject island violations are not judged to be unacceptable. Here, we ask about whether they are acceptable in embedded wh-questions: Abeillé et al. (2020) did not test embedded wh-questions. Our question is whether they will pattern with matrix wh-questions, where subject island violations are unacceptable, or with relative clauses, where they are not. To this end, we designed a 3 (Clause Type) $\times$ 2 (Position) acceptability rating study to test subject island effects in embedded wh-questions. The Position factor had two levels: NP_{subj} involved *wh* phrases extracted from an NP in subject position (hence, a potential subject island violation). NP_{obj} involved *wh* phrases extracted from an object NP, which served as the baseline. The Clause Type factor had three levels: matrix *wh* question, embedded *wh* question, and relative clauses. None of the stimuli contained any form of ellipsis. Sample stimuli in all six conditions are presented in **Table 4**.

	NP_{subj} (potential subject island)	**NP_{obj} (no island, baseline)**
Matrix *wh* question	Which superhero did the costume of amaze Jo at Halloween?	Which superhero did Jo love the costume of at Halloween?
Embedded *wh* question	Megan forgot which superhero the costume of amazed Jo at Halloween.	Megan forgot which superhero Jo loved the costume of at Halloween.
Relative[2] Clause	Bill favors a superhero, which the costume of amazed Jo at Halloween.	Bill favors a superhero, which Jo loved the costume of at Halloween.

Table 4: Sample stimuli for Experiment 3.

[2] Abeillé et al. (2020) found a difference between preposition stranding (as in the example in **Table 4**), and pied piping (e.g., "Bill favors a superhero, of which the costume amazed Jo at Halloween"). Tokens with pied-piping were judged more acceptable, and tokens with preposition stranding did seem to show a subject island effect. Abeillé et al. suggest that this penalty for preposition stranding is not an artefact of extraction out of subjects, but rather, it reflects an orthogonal acceptability penalty for stranding a preposition in a preverbal position. They hypothesize (p.10) that "the difficulty in extracting from a preposition-stranded subject is due to the fact that most cases of P-stranding occur post-verbally in English: NP extractions from nominal subjects are very rare, and hence difficult to process based on

5.1 Method

5.1.1 Materials

Eighteen critical item sets were created and put in a Latin-Square design. Six lists were created, and each participant saw 3 items in each condition. Eighteen filler sentences (half grammatical, half ungrammatical) were recycled from the previous acceptability experiment (Experiment 2). The full set of 36 items was presented in a pseudo-randomized order (no two adjacent critical items of the same condition).

5.1.2 Participants

72 participants were recruited online on MTurk. All self-reported as L1 English speakers. The recruitment procedure was the same as Experiment 2: Participants were master workers located in the US and they received $2 US in exchange for participation. Participants who had completed the previous acceptability experiment (Experiment 2) were excluded from the recruitment pool because a sub-part of the stimuli was recycled and utilized in the current experiment.

5.1.3 Procedure

The experiment was created using the online software PCIbex Farm (Zehr & Schwarz 2018) and conducted on MTurk. Participants read and accepted a disclaimer before starting the experiment. In the experiment, participants were instructed to read English sentences and rate their acceptability on a seven-point scale (1-unacceptable, 7-acceptable). The experiment took approximately 20 minutes to complete.

5.2 Results and discussion

We report the results from all participants, which involve 1296 critical trials. The average rating for the grammatical filler sentences was 6.25 (SD = 1.10), while the ungrammatical fillers received an average rating of 1.71 (SD = 1.07). The findings show that subject conditions were overall rated lower than object conditions (3.78 vs 4.74). The difference in ratings between subject and object conditions was largest in matrix *wh* question conditions (3.74 vs 5.04), suggesting a subject island effect, whereas this difference was smallest in relative clause conditions (3.98 vs 4.44). Importantly, embedded *wh* questions showed a pattern similar to matrix *wh* questions (3.63 vs 4.75, see **Figure 3**).

---

syntactic surprisal (Chaves 2013; Hale 2001; Levy 2008).” For the purposes of our current experiment, we note that pied-piping is not possible in the embedded questions used in the experiment, so we had to use p-stranding across all our conditions and items. Whether the unacceptability we find is due to the subject island, or to an independent ban on p-stranding in a preverbal position, it is repaired by ellipsis.

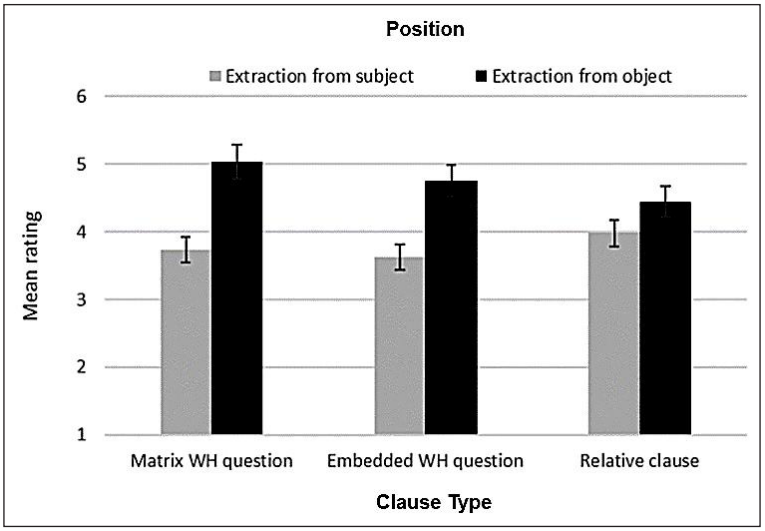

Figure 3: Mean ratings across matrix wh question, embedded wh question, and relative clause conditions for extraction from subject NPs (gray bars) and for extraction from object NPs (black bars). Error bars represent 95% confidence intervals.

Like with Experiment 2, raw acceptability ratings were z-score transformed prior to statistical analysis. Position was sum coded. The three levels of the Clause Type predictor were coded with repeated contrasts: The first coefficient contrasted the acceptability of sentences in the matrix *wh* question condition (coefficient: –2/3) with the embedded *wh* question (coefficient: +1/3 for embedded *wh* question conditions; +1/3 for relative clause conditions). The second coefficient contrasted the embedded *wh* question condition (coefficient: –1/3) with the relative clause condition (coefficient: +2/3 for relative clauses; –1/3 for matrix *wh* questions). We then followed the same statistical procedure as for Experiment 2. **Table 5** shows the results.

EFFECT	*β*	*SE*	*t*	*p*
Position	**0.46**	**0.03**	**14.38**	**<.0001**
Matrix*wh* vs embedded*wh*	**–0.1**	**0.04**	**–2.54**	**.011**
Embedded*wh* vs. RC	0.01	0.04	0.37	.7
Position × Matrix*wh* vs embedded*wh*	–0.1	0.08	–1.35	.18
Position × Embedded*wh* vs. RC	**–0.29**	**0.08**	**–3.75**	**.0002**

Table 5: Results of the linear mixed-effects regression model, Experiment 3. Significant effects are bolded.

Our analysis showed that Position was a significant predictor: Sentences where a wh-phrase is extracted from a subject NP were rated significantly lower compared with extraction from an object (3.78 vs 4.74; $\beta = 0.46$, $SE = 0.03$, $t = 14.38$, $p < .0001$; $BF_{10} = 4562$).[3] There was a significant overall difference between ratings for matrix *wh* questions and embedded *wh* questions (4.39 vs 4.19; $\beta = –0.1$, $SE = 0.04$, $t = –2.54$, $p = .011$); however, the Bayes Factor analysis provided at best weak evidence in support of this difference ($BF_{10} = 1.3$). The overall difference between embedded *wh* questions and relative clauses was not significant ($p = .7$; $BF_{10} = .04$).

However, our key question about whether embedded questions pattern with matrix *wh* questions or with relative clauses with respect to the subject island effect concerns the *interactions*. Importantly, the interaction of position with the comparison between matrix *wh* questions and embedded *wh* questions was *not* significant ($p = .18$), and the Bayes Factor analysis returned a BF_{10} of 0.21, thus providing moderate evidence that matrix and embedded *wh* questions do respond in the same way to the manipulation of position. However, the interaction of position with the comparison between embedded *wh* questions and relative clauses *was* significant ($\beta = –0.3$, $SE = 0.08$, $t = –3.75$, $p = .0002$; $BF_{10} = 91.3$), providing very strong evidence that the subject island effect is different in relative clauses vs embedded *wh* questions. In short, matrix *wh* questions and embedded *wh* questions show a subject island effect, and this is attenuated in relative clauses (cf. Abeillé et al. 2020).

Thus, this experiment shows that embedded *wh* questions pattern with matrix *wh* questions regarding the subject island effect. Recall that the types of sluices we looked at in our Experiment 2 are structurally akin to embedded *wh* questions (e.g., Merchant (2001) and the work following) with an unpronounced embedded clausal complement. The current findings weaken the argument that the nature of subject islands causes the lack of island effect under sluicing because embedding a *wh* question does not attenuate the subject island effect, the way relative clauses do. At the same time, it is still an open question whether non-pronunciation (i.e., ellipsis) repairs island violations under contrast sluicing as claimed in Yoshida et al. (2019). We will revisit this next in our General Discussion.

---

[3] This is true also for relative clauses, which are characterized by Abeillé et al. (2020) as not exhibiting a subject island effect at all. The lower ratings here for extraction out of subjects compared with out of objects thus may reflect a more general penalty (i.e., not tied specifically to extraction) for having an NP with a PP-complement in subject position compared with object position, which we also observed as a main effect in Experiment 2. It may also reflect an orthogonal penalty for preverbal preposition stranding (i.e., as per the stance taken by Abeillé et al.), which we mentioned in footnote 2.

6 General discussion

This study is the first to quantitatively explore how regular and contrast sluicing respond to (subject) islands in English. First, we found distinct patterns in regular sluicing and contrast sluicing concerning the syntax of the elided clause. In sentences compatible with contrast sluicing, there was a notable preference for sentence continuations that were syntactically isomorphic to the antecedent clause. This supports the common structural view that the elided clause in contrast sluicing must be isomorphic with its antecedent. On the other hand, regular sluicing permits and even prefers non-parallel syntactic structures: Experiment 1 found a strong preference for a copular continuation over an isomorphic continuation for sentences compatible with regular sluicing. We interpret this finding as indirect evidence in favor of the evasion approach, which says that island violations can appear to be repaired in sluicing because they can in fact be evaded. Under this approach, the clause that is elided does not actually contain an island (e.g., Barros et al. 2014). If this is correct, a special ellipsis repair mechanism would not be needed to explain the acceptability of regular sluicing with islands (cf., Merchant 2008). However, we acknowledge that our Experiment 1 findings could reveal preferences for pronounced continuations only, and may not be indicative of the unpronounced structure in regular sluicing. Before making any strong conclusions, we next aimed to quantify the island sensitivity of sluicing, particularly contrast sluicing. Contrary to predictions (e.g., Merchant 2008; Barros et al. 2014; Griffiths & Lipták 2014), Experiment 2 found that contrast sluicing is not sensitive to the subject island constraint, while Experiment 3 showed that pronounced versions of these sluicing clauses are. Experiment 3 found that embedded *wh* questions pattern with matrix *wh* questions and against relative clauses (Abeillé et al. 2020) in being sensitive to subject islands. Experiments 2 and 3 together challenge the consensus view in the literature according to which contrastive ellipsis is island-sensitive (see also Barros 2014b; Weir 2014; Yoshida et al. 2019). We suggest that previously reported informal judgments of island sensitivity in contrast sluicing may reflect a general degradation in acceptability compared to regular sluicing, rather than any effect of islands in particular. This general degradation was evident in Experiment 2, where non-island violating contrast sluicing sentences were rated lower in acceptability than their regular sluicing counterparts.

The findings of Experiments 1 and 2/3 seem to point in opposite directions. Experiment 1 suggests that evasion approaches to the island insensitivity of regular sluicing are on the right track. Since contrast sluicing does not allow a non-isomorphic continuation, we would then expect contrast sluicing to be sensitive to islands, as has been reported in the literature. This is not what Experiment 2 showed, however; that experiment found that contrast sluicing is not unacceptable when the presumed extraction takes place from a subject island. Experiment 3 verified that subject islands do constrain embedded question formation. Experiments 2 and 3,

then, indicate that it is something about ellipsis in particular that gets around an island violation, as is proposed by ellipsis repair approaches. Taken together, our results suggest that the evasion approach and the repair approach might both be correct. If what is elided in sluicing is a simpler clause with no island (a copular clause, for instance), then that will certainly evade an island violation in regular sluicing. At the same time, it appears that ellipsis can also repair an island violation, since we found that contrast sluicing is not sensitive to subject islands, but contrast sluicing requires a syntactically isomorphic continuation. Since evasion is not a possibility in contrast sluicing (e.g., Barros et al. 2014), only repair-by-ellipsis could be behind the lack of an island effect.

One might think that allowing both repair by ellipsis and evasion is not the most parsimonious approach. However, once we allow repair by ellipsis, allowing evasion in addition does not require any additional theoretical machinery. So long as a non-isomorphic structure is possible for the elided clause (and it seems to be, if continuations are indicative), then it follows that there will be no island violation, since the structure contains no island. It is the repair-by-ellipsis approach itself that needs additional theoretical machinery. We believe that our results can actually help to reduce this machinery: If our findings are correct and there is no difference between contrast sluicing and regular sluicing in sensitivity to islands, then the different structures and mechanisms that have been proposed in the repair-by-ellipsis approach to account for the presumed contrast can be done away with. There only needs to be one structure and one repair mechanism. For example, Merchant (2008) proposes that in contrast sluicing, the focused correlate (e.g., *Captain America*) can only move as high as VP. In this position, it is not parallel to the wh-phrase in Spec-CP of the elided clause (e.g., *which other superhero*), and so ellipsis of IP is not licensed. This is in contrast with regular sluicing, where an indefinite correlate (e.g., *a superhero*) can apparently move as high as Spec-CP. Given our findings, this unmotivated distinction can be done away with: In both cases, the correlate is capable of moving to (or taking scope at) Spec-CP, providing the needed parallelism for ellipsis licensing.

An alternative to allowing both repair by ellipsis and evasion is to say that there is no evasion, there is only a general repair-by-ellipsis mechanism. On this alternative, pronounced continuations do not give us evidence about the structure of the elided clause in (regular) sluicing. Pronounced continuations are only evidence of the structure of pronounced continuations. We will leave this as an open possibility, although we do think that pronounced continuations are a clue to elided structure (because speakers have to be able to spell out an elided structure when required to).

Finally, non-syntactic approaches to islands, like the discourse-based analysis of Abeillé et al. (2020), will need to explain why ellipsis remedies whatever non-syntactic problem is supposed to give rise to an island effect. For Abeillé et al., this is a discourse clash in the case of

subject islands: wh-phrases are focused, but originating inside the subject makes the wh-phrase backgrounded, and this conflict is what causes unacceptability. In sluicing and contrast sluicing, we cannot see that eliding the clause out of which the wh-phrase has been extracted should change either focus or backgrounding. In fact, the ellipsis is consistent with what is assumed to be the focus structure: The focused phrase is pronounced (the wh-remnant), while the backgrounded material is elided (the rest of the clause). In addition, it is clear that the meaning of the elided clause can be computed (it must be, to be interpreted). We do not see any obvious way to account for the saving effect of ellipsis in this discourse-based approach. It is therefore incumbent upon proponents of this approach, or any other non-syntactic approach to subject islands, to show how their approach can capture the saving effect of ellipsis.

Overall, the recent uptick in research which explores clausal ellipsis phenomena and island phenomena experimentally has helped to clarify the empirical picture and to direct analytical theory. By looking at how regular and contrast sluicing respond to subject islands, our findings have contributed in this same way. We expect that this trajectory should lead to much new exploration of the issues involved in clausal ellipsis and extraction phenomena.

Data availability

The stimuli and results for Experiments 1, 2 and 3, along with Bayesian models results, are publicly accessible through an OSF repository: https://osf.io/e5agt/?view_only=e6ee86c26d1049f7908322c9e5561ad9.

Acknowledgements

We thank Satoshi Tomioka, Rui Chaves and the audience at the Syntax Semantics Lab at the University of Delaware for their insightful comments and discussions. All remaining errors are ours.

Competing interests

The authors have no competing interests to declare.

References

Abeillé, Anne & Hemforth, Barbara & Winckel, Elodie & Gibson, Edward. 2020. Extraction from subjects: Differences in acceptability depend on the discourse function of the construction. *Cognition* 204. 104293. DOI: https://doi.org/10.1016/j.cognition.2020.104293

Anand, Pranav & Hardt, Daniel & McCloskey, James. 2021. The Santa Cruz sluicing data set. *Language* 97(1). e68–e88. DOI: https://dx.doi.org/10.1353/lan.2021.0009

Avetisyan, Serine & Lago, Sol & Vasishth, Shravan. 2020. Does case marking affect agreement attraction in comprehension? *Journal of Memory and Language* 112. 104087. DOI: https://doi.org/10.1016/j.jml.2020.104087

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

Barker, Chris. 2013. Scopability and sluicing. *Linguistics and Philosophy* 36(3). 187–223. DOI: https://doi.org/10.1007/s10988-013-9137-1

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

Barros, Matthew. 2014a. A non-repair approach to island sensitivity in contrastive TP ellipsis. In *Proceedings of the 48th annual meeting of the Chicago Linguistic Society*, 61–75. Chicago, IL: Chicago Linguistics Society.

Barros, Matthew. 2014b. *Sluicing and identity in ellipsis*. New Brunswick, NJ: Rutgers University dissertation.

Barros, Matthew & Elliott, Patrick D. & Thoms, Gary. 2014. There is no island repair. Ms., Rutgers University, University College London, and University of Edinburgh. Retrieved from https://ling.auf.net/lingbuzz/002100

Bates, Douglas & Mächler, Martin & Bolker, Ben & Walker, Steve. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67(1). 1–48. DOI: https://doi.org/10.18637/jss.v067.i01

Bürkner, Paul-Christian. 2017. brms: An R package for Bayesian multilevel models using Stan. *Journal of statistical software* 80. 1–28. DOI: https://doi.org/10.18637/jss.v080.i01

Chaves, Rui P. 2013. An expectation-based account of subject islands and parasitism. *Journal of Linguistics* 49(2). 285–327. DOI: https://doi.org/10.1017/S0022226712000357

Chaves, Rui P. & Dery, Jeruen E. 2019. Frequency effects in Subject Islands. *Journal of Linguistics* 55(3). 475–521. DOI: https://doi.org/10.1017/S0022226718000294

Chomsky, Noam. 1972. Some empirical issues in the theory of transformational grammar. In Peters, Paul S. (ed.), *Goals of linguistic theory*, 63–130. Englewood Cliffs: PrenticeHall Inc.

Chomsky, Noam. 1973. Conditions on transformations. In Anderson, Stephen & Kiparsky, Paul (eds.), *A Festschrift for Morris Halle*, 232–286. Holt, UK: Rinehart and Winston.

Chung, Sandra & Ladusaw, William A. & McCloskey, James. 1995. Sluicing and logical form. *Natural Language Semantics* 3(3). 239–282. DOI: https://doi.org/10.1007/BF01248819

Culicover, Peter W. & Jackendoff, Ray. 1995. Something Else for the Binding Theory. *Linguistic Inquiry* 26. 249–275. DOI: https://doi.org/10.1093/acprof:oso/9780199271092.003.0011

Erteschik-Shir, Nomi. 1973. *On the nature of island constraints*. Cambridge, MA: MIT dissertation.

Ginzburg, Jonathan & Sag, Ivan A. 2000. *Interrogative investigations: The form, meaning, and use of English interrogatives*. Stanford, CA: CSLI Publications.

Goldberg, Adele E. 2006. *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199268511.001.0001

Griffiths, James. 2019. A Q-based approach to clausal ellipsis: Deriving the preposition stranding and island sensitivity generalisations without movement. *Glossa: A Journal of General Linguistics* 4(1). 1–41. DOI: https://doi.org/10.5334/gjgl.653

Griffiths, James & Lipták, Anikó. 2014. Contrast and island sensitivity in clausal ellipsis. *Syntax* 3. 39–52. DOI: https://doi.org/10.1111/synt.12018

Hale, John. 2001. A probabilistic Earley parser as a psycholinguistic model. In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics*. 159–166. Retrieved from https://aclanthology.org/N01-1021.pdf

Huang, C.-T. James. 1982. *Logical relations in Chinese and the theory of grammar*. Cambridge, MA: MIT dissertation.

Jeffreys, Harold. 1998. *The theory of probability*. Oxford University Press.

Kuznetsova, Alexandra & Brockhoff, Per B. & Christensen, Rune Haubo Bojesen. 2017. lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software* 82(13). 1–26. DOI: https://doi.org/10.18637/jss.v082.i13

Lasnik, Howard. 2001. When can you save a structure by destroying it? In *Proceedings of the North East Linguistic Society* 31. 301–320. Georgetown University: GLSA.

Levy, Roger. 2008. Expectation-based syntactic comprehension. *Cognition* 106(3). 1126–1177. DOI: https://doi.org/10.1016/j.cognition.2007.05.006

Merchant, Jason. 2001. *The syntax of silence: Sluicing, islands, and the theory of ellipsis*. Oxford University Press.

Merchant, Jason. 2005. Fragments and ellipsis. *Linguistics and Philosophy* 27. 661–738. DOI: https://doi.org/10.1007/s10988-005-7378-3

Merchant, Jason. 2008. Variable island repair under ellipsis. In Johnson, Kyle (ed.), *Topics in ellipsis*, 132–153. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9780511487033.006

Nykiel, Joanna & Kim, Jong-Bok. 2022. Fragments and structural identity on a direct interpretation approach. *Journal of Linguistics* 58(1). 73–109. DOI: https://doi.org/10.1017/S0022226720000420

Potter, David. 2017. *The Island (In)Sensitivity of Stripping*. Evanston, IL: Northwestern University dissertation.

R Core Team. 2021. R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. https://www.r-project.org/.

Ross, John Robert. 1967. *Constraints on variables in syntax*. Cambridge, MA: MIT dissertation.

Ross, John Robert. 1969. Guess who? In Blinnick, Robert I. & Davison, Alice & Green, Georgia M. & Morgan, Jerry L. (eds.), *Papers from the fifth regional meeting of the Chicago Linguistic Society*, 252–286. Chicago: University of Chicago: Chicago Linguistic Society.

Stepanov, Arthur. 2007. The end of the CED? Minimalism and extraction domains. *Syntax* 10(1). 80–126.

Vicente, Luis. 2008. Syntactic isomorphism and non-isomorphism under ellipsis. Ms., University of California, Santa Cruz. Retrieved from https://www.academia.edu/download/30793950/vicente-isomorphism.pdf

Weir, Andrew. 2014. *Fragments and Clausal Ellipsis*. Amherst, MA: University of Massachusetts Amherst dissertation. DOI: https://doi.org/10.7275/5823750.0

West, Stephen G. & Aiken, Leona S. & Krull, Jennifer L. 1996. Experimental Personality Designs: Analyzing Categorical by Continuous Variable Interactions. *Journal of Personality* 64(1). 1–48. DOI: https://doi.org/10.1111/j.1467-6494.1996.tb00813.x

Wu, Danfeng. 2022. Island violations in stripping constructions. *Glossa: a journal of general linguistics* 7(1). 1–12. DOI: https://doi.org/10.16995/glossa.7895

Yoshida, Masaya & Potter, David & Hunter, Tim. 2019. Condition C reconstruction, clausal ellipsis and island repair, *Natural Language and Linguistic Theory* 37. 1515–1544. DOI: https://doi.org/10.1007/s11049-018-9433-0

Zehr, Jérémy & Schwarz, Florian. 2018. *PennController for Internet Based Experiments (IBEX)*. DOI: https://doi.org/10.17605/OSF.IO/MD832