

A nearly-exhaustive experimental investigation of bridge effects in English

Abstract

In many languages, finite clause-embedding verbs vary in whether they allow wh-dependencies to cross from the embedded to the matrix clause – a phenomenon we call “bridge effects.” Why bridge effects exist has been the subject of much debate; we argue that contributing to the lack of consensus is the relatively small samples of verbs (12-75 for English) previously tested in the literature. To resolve this issue, we report two new data sets of bridge effects covering a nearly-exhaustive sample of 640 English verbs. We use these data sets to address three research questions: Are there bridge effects at all? How well do leading theories of bridge effects explain observed variation across the full range of verbs? And are there new patterns emerging from our data that could lead to a better theory? We ultimately argue in favor of a multivariate approach, drawing upon existing ideas while including a novel morphosyntactic licensing component identified from our data. We also discuss implications for theories of locality and explore how context might affect the acceptability of wh-dependencies.

Keywords

bridge verbs, wh-dependencies, locality constraints, sentence processing, pragmatics, English

Acknowledgments

[Anonymized for review]

(main text: 13,731 words, inclusive of notes, charts and tables, but excluding references)

1. Introduction

Wh-dependencies (also called filler-gap dependencies) appear to be unbounded by distance: a wh-item like *what* can be separated (“extracted”) from its point of semantic interpretation, called its *gap* location by analogy to the position it would occupy in a declarative sentence, by any arbitrary distance, calculated either linearly in number of words or structurally in number of clauses. At the same time, there appear to be non-distance locality constraints on wh-dependencies. The one we focus on in this article is that wh-items can only originate within the complement clause of certain clause-embedding verbs, such as *say* and *think* (1), which are commonly referred to as “bridge verbs” (so named by Erteschik-Shir 1973; first observed by Dean 1967). Extraction from the complement of other verbs, such as *shout* (2), is less acceptable. We will call these between-verb differences in acceptability “bridge effects.”

(1) What did Jo think that Sam said that Kim saw _?

(2) ??What did Jo shout that Sam said that Kim saw _?

As one might expect from over 50 years of research, the space of theories of bridge effects is fairly robust. There are at least three prominent theories that can cover the full range of clause-embedding verbs, each positing a distinct source for bridge effects: information structure-based theories (e.g., Erteschik-Shir 1973, Goldberg 2006, Ambridge & Goldberg 2008, Richter & Chaves 2020), template-based processing theories (e.g., Dąbrowska 2008, 2013; see also Verhagen 2005, 2006), and frequency-based processing theories (e.g., Kothari 2008, Liu et al. 2022). We will review these in detail in Section 2. Each has amassed some amount of experimental evidence in its favor. For example, Ambridge and Goldberg (2008) report an exceptionally strong correlation between backgroundedness and bridge effects (Pearson $r = -.83$, $p = .001$), which supports information structure approaches; Liu et al. (2022) present evidence suggesting there are no bridge effects at all (i.e., no differences in acceptability among verbs) once the frequency of the finite complement clause is taken into account, supporting an extreme version of the frequency-based processing approach; while Richter & Chaves (2020), responding to potential methodological issues in Liu et al.’s study, find limited evidence for a frequency-based approach, arguing instead in favor of an approach based on semantics and information structure. The picture that emerges is of a highly variable empirical landscape that makes theoretical progress difficult.

Recent empirical studies of bridge effects have recognized that the variability in the results is at least partially related to variability in how many verbs were studied. This has led to an increase in sample sizes over time: 8 in Featherston 2004 in German; 12 in Ambridge & Goldberg 2008; 24 and 48 in Liu et al. 2022; 75 (experimental) and 136 (corpus) in Richter & Chaves 2020. Given that the results appear to change with sample sizes, the logical conclusion is that the field could benefit from testing a comprehensive set of verbs. Though creating the data set would be resource-intensive, it would both license more accurate evaluations of existing theories and open new pathways for theory construction. It would also act as a benchmark data set, thereby relieving each study of the need to invest resources in collecting yet another sample of verbs. To that end, we collected two extremely large-scale data sets of bridge effects for 640 finite clause-embedding verbs in English – a nearly exhaustive list of such verbs in English (see Appendix A). The first measured acceptability with sentences presented in isolation; the second with preceding (supportive) context. So that we have good estimates, each experiment targeted responses from 60 participants per verb, requiring the recruitment of over 9,000 participants on Amazon Mechanical Turk (Section 3). We also compiled measures for each verb that are relevant to theories of bridge effects (Section 4). For information-structure theories, this entailed a third large-scale experiment to collect backgroundedness judgments using over 5,000 additional participants. For template-

based processing theories, we collected twelve measures of semantic similarity from the natural language processing literature. And for frequency-based processing theories, we collected two measures of frequency from COCA. We [will make] our data sets publicly available on our websites for researchers to use in their own studies of bridge effects.

As a first set of studies with these new, nearly exhaustive data sets, this article addresses three research questions.

Our first question is empirical: Do bridge effects exist? Section 5 evaluates this question on our full set of verbs, responding to two claims in the literature. First, Liu et al. (2022) argue that once the frequency of the verb co-occurring with finite clauses is accounted for, bridge effects disappear (i.e., there is a main effect of frequency and a main effect of extraction on acceptability, but no interaction). Second, information structure accounts have reported that supportive context can improve long-distance wh-extraction for certain non-bridge verbs, inviting the prediction that context can reduce or even eliminate bridge effects. Using the analysis approach adopted in Liu et al., we find that bridge effects exist even after accounting for each of the main predictors from each of the three prominent theories: frequency (contra Liu et al.), backgroundedness, and semantic similarity. With regard to context, we find that context does reduce the penalty for wh-extraction, but the effect is relatively small on average. Crucially, context neither eliminates nor reduces the between-verb differences in penalties that characterize bridge effects, contrary to what one might expect from information structure accounts.

Our second question is the central theoretical debate in the literature: What is the source of bridge effects? Section 6 presents confirmatory analyses to evaluate predictions of the three theories, using the analysis approach adopted in Ambridge & Goldberg 2008: simple linear regressions between the predictor variables for each theory and our acceptability data. Anticipating our results, we find that the information structure-based theory performs best overall, but that (subjectively) none of the fits are particularly strong. This provides a window for understanding the seemingly incompatible results across studies – the relatively poor fits are more susceptible to sampling error based on the number of verbs tested. Our results further suggest that research on bridge effects could benefit by considering a wider range of theories beyond these three.

Therefore, our third question is exploratory: Are there new patterns visible in our data sets that could lead to a better theory? Section 7 reports just such a pattern: the penalty for long-distance extraction appears to correlate to a fair degree with whether a verb also selects for nonfinite complement clauses. As far as we know, this morphosyntactic property has not been noted before (and indeed, syntactic factors are underexplored in research on bridge effects). We first discuss what this property may be, building on work by Wurmbrand (2019). Then, building on suggestions in Erteschik-Shir 1973, Richter & Chaves 2020, and Chaves & Putnam 2020 that multiple factors may be required to explain bridge effects, we construct a multivariate theory that combines this new morphosyntactic property with the best predictors from existing theories. We show that this theory delivers better empirical coverage, even after controlling for complexity. This finding affirms the value of a multivariate approach and provides a novel argument for the view that morphosyntax can play an important role in bridge effects.

Finally, in Section 8, we consider two implications of our results for broader questions beyond bridge effects. We discuss results in Section 6 showing that bridge effects have smaller effect sizes (about .3 z-units) compared with island effects as a potentially relevant fact for efforts to unify bridge effects and island effects within a single theory. We also build on an observation in Section 6 that model fits are higher for sentences presented with context as a way to begin to explore the effect of context on the acceptability of long-distance dependencies.

2. Review of existing accounts

2.1. Defining long-distance penalties and bridge effects

Before reviewing the literature, it will be helpful to have a precise definition of bridge effects that follows both the theoretical and the experimental literature. The first step is to define the long-distance penalty that occurs for extraction (for all verbs). We follow Ambridge & Goldberg (2008) in defining the long-distance penalty for a verb as the acceptability difference between a wh-question with extraction from the matrix clause (short dependency) and a wh-question with extraction from the complement clause (long dependency), as shown in (3).

- (3) Long-distance penalty = short dependency rating – long dependency rating
- a. ??What did Jo shout that Sam saw _? *long dependency*
 - b. Who _ shouted that Sam saw the movie? *short dependency*

The direction of this subtraction means that penalties will be positive when long-distance extraction is worse than short. We note that this matches Ambridge & Goldberg's discussion, but not the general norm in the experimental literature to perform subtractions in the other direction (experimental condition – control condition). The second step is to define “bridge effects” as differences in penalties between verbs – intuitively, the penalty for *shout* (a non-bridge) will be a larger positive number, and the penalty for *say* (a bridge) will be a smaller positive number. Again, this is consistent with how bridge effects have been defined in the empirical and theoretical literatures. However, we note that some discussions of bridge effects appear to focus exclusively on the variation in the absolute acceptability of long-distance extraction (i.e., just one condition). Taken at face value, this would lead to a confound with the general acceptability of clause-embedding for each verb (as Ambridge & Goldberg 2008 note). We suspect that these discussions are simply using a shorthand – they are assuming that the acceptability of clause-embedding varies less than the acceptability of long-distance extraction, so one can focus exclusively on the interesting condition. But we must include both to quantitatively assess the theories.

2.2. Information structure-based theories

We begin our review by considering one of the first comprehensive theories of bridge effects: Erteschik-Shir's (1973) information structure theory (see also Ambridge & Goldberg 2008). The core of Erteschik-Shir's proposal, intended to cover both bridge effects and island effects, is stated in (4).

- (4) Extraction can occur out of constituents that can be considered dominant in some context (Erteschik-Shir 1973:27), where “dominant” is best understood as “natural to comment on” (Erteschik-Shir *ibid.*:16) or as “focusable,” allowing “the speaker ... to draw attention of the hearer” to the constituent (2017:7).

Erteschik-Shir (1973) provides a number of diagnostics for dominance/focusability. One such diagnostic is the “lie test” (suggested by J. R. Ross). This test shows that, in (5), both the matrix clause and the complement clause of *think* can be dominant, because their propositions can be challenged as lies. In contrast, in (6), only the matrix clause is dominant, because it is infelicitous to challenge the complement clause of *shout* or *know* as a lie.

- (5) Fred thinks that Mary won.
- a. That's a lie, he doesn't think Mary won.
 - b. That's a lie, she didn't win.
- (6) Fred shouted/knew that Mary won.
- a. That's a lie, he didn't shout/know that Mary won.
 - b. ??That's a lie, she didn't win.

Ambridge & Goldberg (2008:364-366; also Goldberg 2006) reframe Erteschik-Shir's notion of "dominance" (or "focusability") in terms of "backgroundedness," and propose a constraint that prohibits gaps from appearing inside backgrounded constituents. Ambridge & Goldberg support this proposal with a study testing 12 English verbs. For each verb, they estimated penalty scores with an acceptability judgment survey. They also measured how much the verb "backgrounds" its complement clause, using a negation test (7), which is a variant of the lie test. The negation test works as follows: if sentential negation negates the proposition represented by a complement clause, the clause is focused; if sentential negation fails to negate the complement clause, the clause is backgrounded. Ambridge & Goldberg found a strong correlation between backgroundedness and penalty scores (Pearson $r=-.83$, $p=.001$; see also Dąbrowska 2013 for a replication with 16 verbs).

- (7) a. *Maria didn't know that Ian liked the cake.*
 Does not imply *Ian didn't like the cake.* (know backgrounds its clause)
 b. *Maria didn't think that Ian liked the cake.*
 More likely to imply *Ian didn't like the cake.* (think does not)

However, an analysis of only 12 verbs suffers from power issues, which Ambridge & Goldberg themselves note (2008:375). Underpowered studies with statistically significant results tend to overestimate the effect size (e.g. Vasisht et al. 2018). Our study tackles this power issue directly by testing a nearly exhaustive set of 640 verbs.

Information structure-based theories have also claimed that supportive context can make non-bridge verbs salient and their complement clause dominant/focusable, thereby improving the acceptability of wh-extraction from the clause (Erteschik-Shir 1973, 2017, Ambridge & Goldberg 2008, Müller 2015, Chaves & Putnam 2020; cf. Kothari 2008 for reading time experiments). The weak version of this claim is that context will reduce penalties by the same extent for all verbs, leaving bridge effects – i.e., the between-verb differences in penalties – unchanged. A stronger version of this claim is that context will improve long-distance wh-extraction for non-bridge verbs more than bridge verbs, thereby reducing or eliminating bridge effects. It is not always clear which version of the claim is endorsed in a given paper, but we note that the general discussion in the field tends to implicitly assume the stronger claim. Our two nearly-exhaustive data sets, without and with context, allow us to test both versions of the claim.

2.3. Template-based processing theory

A second theory derives bridge effects from difficulty experienced during template-based processing. Dąbrowska (2008, 2013, etc.; see also Verhagen 2005, 2006) argue that the processing of questions with long-distance wh-dependencies (henceforth, "long-distance wh-questions") involves the use of lexical templates like "WH *do you think* S-GAP" or "WH *do you say* S-GAP", where WH and S-GAP are respectively variables for a wh-phrase and a finite clause with a gap. Lexical templates are pre-assembled "lexical formulas" based on the most frequent long-distance wh-questions, which feature second person subjects and a verb like *say* or *think*. Templates free speakers from having to build the representation of a wh-question from scratch. Instead, speakers can insert appropriate phrases into the WH and S-GAP variables in a template.

To process a long-distance wh-question with a verb other than *say* or *think*, one must further alter the template by replacing the verb. As Ambridge & Goldberg suggest, bridge effects might reflect how easy replacing the verb is, which in turn depends on how semantically similar the replacement verb is to *say* or *think*. In this view, long-distance wh-extraction incurs a small penalty

for *claim* because *claim* and *say* are semantically similar, and so replacing *say* with *claim* is easy. *Shout*, on the other hand, is associated with a higher penalty because *shout* is less similar to *say*.¹

This account differs from the information structure theory in two ways. First, bridge effects here are language processing artifacts. Second, this account is not as complete a theory, as it does not explain why *say* and *think* can appear with long-distance wh-dependencies in the first place.

We note that Ambridge & Goldberg (2008:380-382) found little evidence for this theory: bridge effects showed no correlation with similarity measures derived from a judgment survey and a Latent Semantic Analysis calculator (LSA; Deerwester et al. 1990) for their set of 12 verbs. However, one could wonder if their null result might reflect issues with their sample of verbs and/or LSA data set. We address these potential concerns by re-evaluating this theory through our much larger set of verbs and also with four measures of semantic similarity (see Section 4.2).

2.4. Frequency-based processing theory

The original frequency-based processing theory posits that verbs that appear more frequently with a finite complement clause will have smaller penalties for extraction from the clause. This builds on the idea that frequency is straightforwardly correlated with processing difficulty and/or acceptability (e.g. Hale 2001, Levy 2008; but see Sprouse et al. 2018, White & Rawlins 2020). Some data in favor of such a view can be found in Kothari 2008, which showed that frame frequency (the frequency that a verb appears with a finite clause) and a verb's bias for a finite clause (i.e. the conditional probability of such a clause given a verb) are both correlated with the absolute acceptability of long-distance wh-questions.

In more recent work, however, Liu et al. (2022) have taken the frequency-based processing theory a step further, and argued that there are in fact no bridge effects once frame frequency is taken into account. Their experiments compare acceptability for a condition with extraction to a condition without extraction for 24 and 48 verbs with varying frame frequencies: they find a main effect of extraction (i.e., a penalty for all long-distance extraction) and a main effect of frame frequency (frequency being correlated with acceptability), but, crucially, no interaction between these two factors. This means that they observe no bridge effects – no meaningful differences in the size of the penalties between verbs.

There are several methodological reasons to follow up on the Liu et al. result. For one, as Richter & Chaves (2020) note, the lack of statistical interaction might reflect participant fatigue since Liu et al.'s experiments were designed to be very long (e.g. 288 items in their Experiment 2) and uniform (all items were critical items, with no filler items). Liu et al.'s experiments also only tested samples of 24 and 48 verbs – larger than Ambridge & Goldberg but still relatively small compared to the full range of clause-embedding verbs in English. Richter & Chaves (2020) begin to address this with an experimental study of 75 verbs, ultimately concluding against Liu et al. Our study takes this to the logical conclusion by testing a nearly exhaustive set of verbs (640) using much shorter experiments (31 items) with a 2:1 ratio of (pre-tested) fillers to target items.

3. Quantifying bridge effects using acceptability judgment experiments

In this section, we describe how we obtained quantitative measures of bridge effects, by compiling a set of 640 clause-embedding verbs and collecting the acceptability of wh-questions containing these verbs, in isolation and after a context sentence. In the next section (Section 4), we describe how we compiled predictors for each of the three theories reviewed above. Readers should feel free to skip ahead to the results (starting from Section 5).

¹ This formulation departs from Dąbrowska 2008, 2013. For Dąbrowska, verb similarity affects the absolute acceptability of long-distance wh-questions, but not necessarily bridge effects (see Section 2.1).

3.1. Compiling a nearly exhaustive set of finite clause-embedding verbs

We assembled 640 verbs (Appendix A) from Anand, Grimshaw & Hacquard 2019; Levin 1993; and the MegaVeridicality data set (White & Rawlins 2018). We focused on verbs whose active voice forms must assign thematic roles to a subject. This criterion includes verbs like *say*, *shout*, and other verbs canonically used to illustrate bridge effects, while excluding raising verbs like *seem* and psych-verbs like *surprise*, which can appear in the active voice with an expletive *it* subject. From the resulting set of 641 verbs, we excluded the verb *animadvert* because it is so rare as to be absent from the one-billion-word Corpus of Contemporary American English (COCA; Davies 2020), leaving 640 verbs.

3.2. Acceptability judgment experiment 1: judgments in isolation (no context)

3.2.1. Constructing wh-question items

We first created a set of 10 different frames, each with a distinct combination of nouns and verbs. An example is shown in (8) with the verb *tell*. These frames were intended to be semantically compatible with a majority of the 640 clause-embedding verbs, so that any acceptability variation can be reasonably attributed to the verb being incompatible with long-distance wh-dependencies and not to plausibility.

(8) The party leader told the vice-president that the senators would endorse the governor.

For a subset of 286 verbs, however, we judged that certain frames were unsuitable for specific lexical semantic reasons. For instance, *broadcast* and *editorialize* prefer subjects denoting media organizations, like *the TV station*. We sorted these verbs into 52 smaller classes, based on our intuitions, and adapted the frames accordingly.

The tense and modality of the complement clause also varied for similar reasons. To maximize comparability, we made *would* the default tense/modal marker, as in (8), because some verbs have predictive semantics (e.g. *predict*, *expect*) and are most felicitous with future modals. But we changed the tense/modal marker in the complement clause where required by the verb.

The 10 frames (and modifications) lead to 10 lexically matching pairs of short (9) and long (10) wh-questions for each clause-embedding verb, i.e. two conditions for each verb, consistent with our operationalization of bridge effects (Section 2.1). Clause-embedding verbs appeared in the past tense, like *told* in (8), with four exceptions: we judged that *stand*, *bear*, and *forgive* should co-occur idiomatically with the modal *couldn't*, while *care* and *mind*, being negative polarity items, should co-occur with the auxiliary *didn't*. As for the wh-word, this was chosen to match the corresponding NP in each frame: if the NP denoted human entities or groups, the wh-word was *who*, otherwise *what*.

(9) Who _ told the vice-president that the senators would endorse the governor?

(10) Who did the party leader tell the vice-president that the senators would endorse _?

The full set of experimental items will be made available along with the acceptability ratings data set.

3.2.2. List creation

To avoid fatigue effects, we made the acceptability judgment surveys short: a total of 31 sentences, consisting of eight target sentences that varied by list and 23 filler sentences that were identical across lists. The eight target sentences consisted of four clause-embedding verbs, each appearing in both short and long conditions, so that we could calculate a within-participant penalty score per verb. We split our 640 verbs into four bins of 160 verbs based on how frequently they occur with complementizer *that* in COCA. We then formed 160 quadruplets to be tested together in a single survey by randomly sampling one verb from each bin, so each quadruplet has a mix of

verb frequencies. We also made sure that each quadruplet had a mix of semantic types (e.g., never all manner-of-speaking).

The 10 frames for each verb were then distributed in a Latin Square design, so that participants would never see the same frame across verbs or conditions in their list. Doing so yields 10 different lists for each of the 160 quadruplets, for a total of 1,600 unique lists.

3.2.3. Filler composition

Our 23 fillers were based on sentences from Sprouse, Schütze & Almeida 2013 that have well-established acceptability ratings and are known to span the range of possible ratings in a 7-point acceptability judgment task. Our fillers were intended to introduce variability in the items, to combat fatigue and boredom, as well as to encourage participants to use the full range of the acceptability scale. Nine of the fillers appeared in a fixed order at the start of each survey; the first seven each have an expected mean of 1 through 7, while the eighth and ninth have an expected mean of 1 and 7. The remaining 14 fillers were composed of two items each with expected ratings of 1 through 7. These 14 fillers and eight target sentences were presented in a random order.

3.2.4. A synonym post-test for determining verb familiarity

Many of our 640 verbs appear infrequently, like *grok* and *expostulate*. We implemented a post-test at the end of each acceptability rating survey to check for participants' familiarity with the verbs. Each of the four verbs surveyed appeared in descending order based on frequency, together with an example declarative sentence formed by one of the 10 frames for that verb. Participants were to identify the synonym, i.e. pick a verb or phrase "closest in meaning," from a set of four choices presented in a random order: a close synonym, an antonym, and two other semantically unrelated verbs. We use this post-test to eliminate both uncooperative participants and trials in which the participant does not know the meaning of the verb (see section 3.2.7).

3.2.5. Task and presentation

Stimuli were presented using the Ibex experiment platform (Drummond 2012). Participants were instructed to rate sentences on a 7-point scale based on whether they think a native speaker of English could say these sentences in a conversation. Participants were instructed to ignore prescriptivist rules. A rating of 1 indicated that the sentence was "very bad" and a rating of 7 "very good." To anchor the scale, three example sentences with suggested ratings of 1, 4, and 7 were embedded within the instructions. These were also taken from Sprouse et al. 2013 with means of 1, 4, and 7. Sentences were presented one per screen. Participants could take as much time as they wanted to respond, but could not go back to previous sentences.

3.2.6. Participants

We recruited participants in two stages, first directly on Amazon Mechanical Turk (AMT) and then through CloudResearch, a platform that allows researchers to target a pool of more reliable AMT participants. All participants were self-reported native speakers of American English, above 18 years of age, and based in the United States. Participants recruited directly on AMT also had to have completed more than 500 tasks and received an approval rating of at least 95% on their previous tasks. Each participant received US\$1 for completing a survey; this was based on an hourly rate of US\$12/hour and our estimate that a survey, with 31 items total, should take at most 5 minutes to complete.

For each survey, we planned to recruit 60 participants (30 directly on AMT and 30 through CloudResearch), in order to collect 60 responses per verb per condition, although we sometimes inadvertently collected more responses due to the mechanics of Ibex and AMT. We tracked participants so that each participant only completed one survey. We also tried as far as possible to minimize overlaps in the AMT and CloudResearch participant pools. Of the 9,219 unique

participants recruited, only 805 participants (8.7%), completed two surveys, once as part of the AMT pool and once as part of the CloudResearch pool. Although the filler items were identical across the two surveys for those participants, the two surveys were run almost two years apart, so the participants were unlikely to remember the sentences.

3.2.7. Data analysis and outlier detection

After data collection, we identified four sentences that were incorrectly presented, and removed those four sentences from analysis. All ratings were z-score transformed by participant to eliminate common forms of scale bias. We only included participants if they were native speakers, read the sentences carefully, responded accurately to the fillers, and knew the verbs, operationalized as follows:

1. They answered yes to two language questions: that they lived in the United States from birth to at least age 13 and that their parents spoke English to them at home.
2. Their median response time to each sentence is at least 2.5 seconds.
3. They responded to at least 12 out of 14 fillers with a rating that is within 2 standard deviations of the mean for that filler.
4. They gave at least three correct responses out of a maximum of four to the synonym post-test.

We included individual trials if they were read carefully and the participant knew the verb, operationalized as:

1. The trial response time was at least 2.5 seconds.
2. The synonym of the verb in the trial was correctly identified in the post-test.

We further excluded a subset of verbs from analysis. We were deliberately liberal in compiling the 640 clause-embedding verbs. Consequently, for certain verbs, clausal complements might be ungrammatical for many native speakers. To identify these verbs, we calculated a mean z-score for each verb in the short wh-dependency condition, where there is no gap in the complement clause. We excluded the 150 verbs with a negative mean z-score for this condition, i.e. below the grand mean of all items in the surveys (designed to have a mean rating near 4, the midpoint of the scale). In addition, for six other verbs, mean short wh-dependency ratings were lower than mean long wh-dependency ratings. This pattern is anomalous under all of the theories under consideration here. We also excluded these verbs out of an abundance of caution. Altogether, these criteria yield a total of 41,958 responses (20,979 pairs of ratings for short and long conditions) for 484 verbs for analysis.²

3.3. Acceptability judgment experiment 2: with prior context

Bridge effects are typically illustrated by presenting wh-questions in isolation. Similarly, as far as we know, all previous experimental studies, except one (Kothari 2008), have tested them in isolation. However, as mentioned in Section 2.2, information structure accounts have claimed that supportive context can impact extraction and possibly penalties and bridge effects. We therefore believe it is important to also test the verbs with prior context.

To our knowledge, there are no detailed theories of how supportive context affects penalties. However, previous work has reported that a dialogue format can improve wh-extraction from complements of non-bridge verbs or island structures (Erteschik-Shir 1973, Chaves & Putnam 2020; see Ambridge et al. 2015:e120 for islands). Here, we adapt the dialogue format used by Ambridge et al. Although the degree of improvement is disputed (see Pérez-Leroux &

² We note that including all 640 verbs in our analyses yields quantitatively similar results to the ones we report below, so these exclusions have no substantive effect. But we believe these exclusions provide the cleanest data set to explore theories of bridge effects.

Kahnemuyipour 2014), this dialogue is one of the few concrete proposals in the locality literature, and therefore seemed like a reasonable candidate for a first systematic exploration of context effects.

3.3.1. Materials

All target items were presented as a dialogue between two individuals, A and B, as in (11) and (12).

- (11) A: Someone thought that the duchess would invite the arrogant knight.
B: Really? Who thought that the duchess would invite the arrogant knight?
- (12) A: The princess thought that the duchess would invite a certain person.
B: Really? Who did the princess think that the duchess would invite?

A's utterance was always an assertion in which a clause-embedding verb appears with a complement clause with no gap. B's utterance was always a wh-question responding to A's assertion, which justifies B's use of the verb and the complement clause. We added *Really?* to signal that B is responding to A's utterance. Participants were instructed to judge only the last sentence in B's utterance, i.e., the question, which was underlined. Before starting, participants saw three similarly-formatted dialogue examples, with suggested ratings.

We constructed these items with the same procedure described in Section 3.2. For short wh-dependency items, A's utterance always featured *someone* or *something* in a subject position, corresponding to *who* or *what* in B's response. For long wh-dependency items, A's utterance always featured *a certain person/thing* in the object position, corresponding to *who/what* in B's response. We did not use *someone/something* because we judged that they would be likelier to receive a nonspecific reading, making it odd to question them.

For fillers, we reused the items described in Section 3.2.3, creating suitable declarative sentences for A as context.

To ensure that participants used the context provided to judge target sentences, we added four catch trials (13)-(16). The target sentences contained the presupposition trigger *either*, licensed if the context sentence was negated, or *too*, licensed if the context sentence was in the affirmative. Participants who read the entire dialogue should notice that (15) and (16) are infelicitous, due to presupposition failure, and give lower ratings to the target sentences.

- (13) A: The carpenter did not repair the table.
B: The apprentice did not repair the table either.
- (14) A: The diver went to the pool.
B: The swimmer went to the pool, too.
- (15) A: The boys ate the broccoli.
B: #The girls did not eat the broccoli either.
- (16) A: The guide did not board the bus.
B: #The tourists boarded the bus, too.

A synonym post-test (see Section 3.2.4) was also included at the end of each survey.

3.3.2. Participants

Participant recruitment proceeded as described in Section 3.2.6; again, we targeted recruiting 60 participants per survey. Participants received \$1.20 for completing a survey. Most participants completed only one survey; only 703 (7.7% of 9,156 unique participants) completed the surveys twice (again, at least two years apart).

3.3.3. Data analysis and outlier detection

For data quality purposes, we applied the same inclusion criteria to the data set, except that:

1. The response time threshold was raised to 3 seconds, since items were longer.

2. We analyzed ratings for the four catch trials to ensure that participants were basing their judgments on the dialogues. Participants must have mean z-scored ratings for the felicitous items that are at least 0.5 units higher than those for the infelicitous items.

So that we can compare ratings for wh-questions with and without prior context, we filtered for the same 484 verbs, yielding a total of 21,732 responses. This number is only about half of the number of responses for the wh-questions presented without context, partly because many participants failed the new catch trial criterion, which requires a difference of 0.5 z-units. Relaxing this criterion to any positive difference only slightly increases the number of included responses (to ~26,000), so we have opted to keep the stricter criterion in place.

4. Obtaining predictors of bridge effects

4.1. Information structure theories: the negation test for backgroundedness

In information structure theories, the key predictor is how much a verb's complement clause is "dominant/focusable" or the inverse, "backgrounded." For our purposes, we assume that this can be quantified using Ambridge & Goldberg's (2008) negation test. For compatibility with their experiment, we will label this measure "backgroundedness."

4.1.1. Task

We collected judgments for the negation test over the internet using Ibex. We pair a sentence in which a clause-embedding verb was negated with another sentence formed by the verb's complement clause, as in (17). Participants were instructed to decide whether the second sentence was true or false, using only information from the first sentence and not any real-world facts. The second sentence was underlined, to make it clear that this was the sentence to judge.

(17) The princess didn't think that the duchess would invite the arrogant knight.

The duchess will invite the arrogant knight.

To the extent that participants judged that the second sentence was true, the verb backgrounds the complement clause. We provided a third option, "Not enough information," in case participants found it difficult to make True/False judgments. Before starting the experiment, participants saw three example items intended to elicit a True response, a False response, and a "Not enough information" response, in that order.

We note that our design departs from prior work. Both Ambridge & Goldberg and Liu et al. used a Likert scale, where the rating indicates how true participants felt the second sentence was, in light of the first sentence. In pilot testing, we noticed that individuals were unlikely to use the full range of a Likert scale, nor use it in a way that corresponds naturally to an ordinal scale: with a 7-point scale, testers reported often using only three options, typically 1 (False), 7 (True) and 4; testers reported using 4 to indicate uncertainty, and not because they felt that the sentence was midway between being True and False. Therefore we judged that a three-way True/False/Not enough information format would capture participant intuitions more transparently. Crucially, it still yields a gradient measure in the sense that each verb could have a different proportion of responses out of the 60 participants per verb (similar to the way Liu et al. use binary acceptability judgments for their experiments).

4.1.2. List creation

Surveys were 27 items long, with eight clause-embedding verbs per survey and 19 fillers. We divided the 640 verbs into eight bins, based on how often a verb occurs with complementizer *that* in COCA. We formed 80 sets of eight verbs by randomly sampling, without replacement, from each bin.

4.1.3. Materials

Verbs and frames were identical to those used for the acceptability judgment task. For the first sentence, all clause-embedding verbs were negated with *didn't*. We applied this negation to the verbs *stand*, *bear*, and *forgive* for consistency, even though doing so produced ungrammatical sentences for *stand* and *bear*, which co-occur idiomatically with *couldn't*, and potentially changes the meaning of the sentence for *forgive*. For the second sentence, we reused the complement clause of the first sentence, except for verbs like *prefer* or *ask*, which have subjunctive complement clauses. For these verbs, we added the modal *should* to the second sentences.

As was the case for the acceptability judgment surveys, we included a synonym post-test for the eight verbs tested in each backgroundedness survey.

4.1.4. Fillers

The 19 filler items each consisted of two sentences. Five fillers were intended to elicit a True response (“True fillers”): The second sentence corresponded to a presupposed clausal subject or adjunct in the first sentence. Eight fillers were intended to elicit a False response (“False fillers”); the first sentence in these fillers featured some kind of negation, while the second sentence was the affirmative variant. The remaining six filler items were likely to elicit a “Not enough information” response: the first sentence in these fillers typically contained a modal adverb or auxiliary, while the second sentence was the affirmative variant. We note that we introduced more False fillers than True fillers. This was an error relative to our initial intentions (an equal number of False and True fillers), but in retrospect, doing so is not unwelcome. Many clause-embedding verbs are factive or invite a factive reading, meaning test items were more likely to be judged True. The extra False fillers may have serendipitously led to a better balance of responses within the experiment.

4.1.5. Participants

Participant recruitment proceeded as described in Section 3, targeting 60 participants per survey, as before. Each of the 5,069 unique participants completed only one survey. Participants received US\$1.20 for completing a survey; this payment rate was based on a US\$12/hour rate and our estimate that each survey, consisting of 27 items and a relatively difficult judgment task, might take slightly over 5 minutes to complete.

4.1.6. Data analysis

For data quality purposes, we imposed criteria that were identical to the ones in Section 3.2.7, except that participants must have given (i) 6 or more correct responses to the synonym post-test and (ii) 9 or more correct responses to the 13 “True” and “False” fillers (i.e. a True response to the True fillers and a False response to the False fillers).

After data collection, we identified one sentence that was incorrectly presented and removed it from analysis. We also excluded all responses for sentences featuring *stand*, *bear*, and *forgive*, because these verbs were negated with *didn't*, but idiomatically co-occur with *couldn't*.

We calculated two related backgroundedness measures. The first measure, %True, was defined as the percentage of True responses for each verb. The second measure, %Not-False, used a more liberal definition: the percentage of responses that were either True or “Not enough information”. The higher these measures, the more backgrounded the complement clause is.

4.2. Template-based processing theories: similarity with *say* and *think*

This theory assumes that the processing of long-distance wh-questions relies on lexical templates featuring *say* and *think*. Bridge effects reflect the difficulty of replacing *say* and *think* in the templates, which depends on how semantically similar the replacement verb is to these two verbs: the greater the similarity, the smaller the penalty.

Ambridge & Goldberg argued against this hypothesis by showing that penalty scores are not significantly correlated with semantic similarity measures derived from Latent Semantic Analysis (LSA). We expand on this work by computing similarity measures using four different data sets from the natural language processing literature:

1. **LSA/Wikipedia** (Ștefănescu et al. 2014). These are word embeddings derived by applying LSA on a 2013 version of English Wikipedia. We chose this data set over the LSA calculator used by Ambridge & Goldberg because this Wikipedia-derived data set is more recent.
2. **Global Vectors (GloVe)/Wikipedia** (Fares et al. 2017). These are word embeddings created by applying GloVe (Pennington et al. 2014), an unsupervised learning algorithm, on a 2017 version of English Wikipedia.
3. **GloVe/Gigaword** (Fares et al. 2017). These are word embeddings created by applying GloVe on the 5th edition of the 4-billion-word Gigaword corpus, compiled from English news outlets (Parker et al. 2011). This provided an alternative to the Wikipedia-based measures above.
4. **WordNet** (Fellbaum 1998). WordNet provides a hierarchy of senses (“synset”) for English words. For each verb, we used WordNet’s definitions and example sentences, as included in Python’s NLTK package (Bird et al. 2009), to identify the synset that most closely corresponds to how the verb is used in the items in the acceptability judgment tasks.

For each data set, we calculated similarity scores for each verb relative to three anchors: similarity with *say*, similarity with *think*, and a hybrid score that chooses whichever score is greater between *say* and *think* (following Ambridge & Goldberg 2008). We define similarity as cosine similarity for the first three data sets and as NLTK’s path similarity measure for the WordNet data. Path similarity tracks the distance between a verb and *think* (or *say*) within WordNet’s hierarchy.

4.3. Frequency-based processing difficulty: frame frequency and verb bias in COCA

To investigate frequency theories, we calculate two measures from COCA.

1. **Frame frequency**. This is how frequently a verb appears with a finite complement clause. Following Liu et al. (2022), we approximate this with how frequently the verb’s lemma immediately precedes the complementizer *that* in COCA.
2. **Verb bias**. This measure is derived by dividing the raw frame frequency by verb frequency, following Richter & Chaves (2020).

We base our estimates on COCA because it is the largest American English corpus, so estimates would align better with our participants’ linguistic experience as American English speakers. Furthermore, COCA is lemmatized and tagged for parts of speech, making it straightforward to identify a verb and a complementizer. However, because COCA does not mark the null complementizer, our frequency values are necessarily underestimates. This issue predominantly affects high-frequency verbs appearing in informal registers (e.g. *say* but not *testify*); which are verbs that most often appear with a null complementizer (e.g. Biber 1999).

5. Do bridge effects exist?

The first question we address with our data sets is whether bridge effects exist in both our “no context” and “with context” data sets. We do so for two reasons. First, recall that Liu et al. (2022) concluded, from their survey of 24 and 48 English verbs, that the acceptability of sentences with clause-embedding verbs presented without context can be modeled in terms of only two main effects (namely, presence of a long wh-dependency and frame frequency) and without an interaction effect, thus effectively denying the existence of bridge effects. Second, one possible

prediction from information structure accounts is that bridge effects might decrease or even disappear in the presence of supportive context.

5.1. Are there differences in penalties between verbs?

Given that we operationalized bridge effects as differences in penalty scores between verbs, a direct test of whether bridge effects exist is to see whether penalties for each verb actually vary. Figure 1 shows substantial variation in both “no context” and “with context” mean penalty scores between the 484 verbs analyzed.

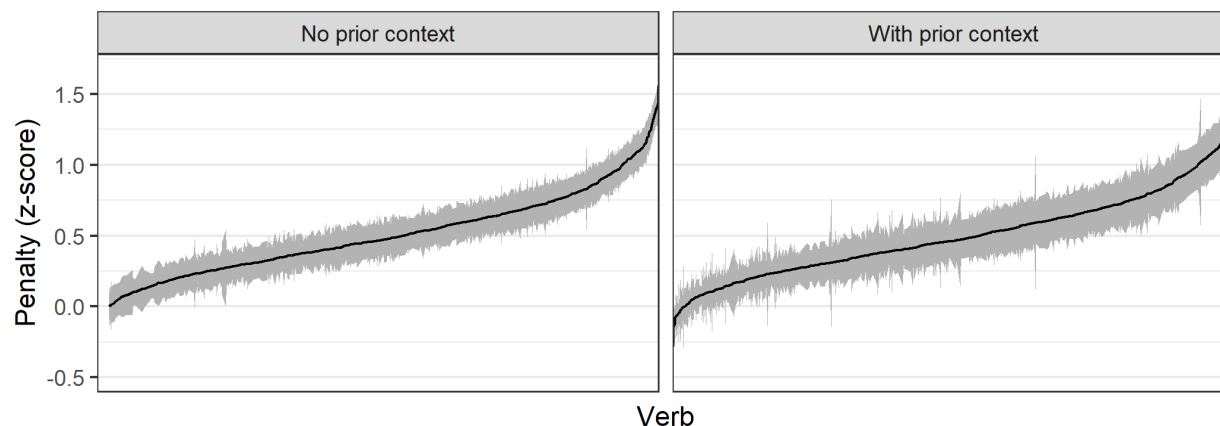


Figure 1: Penalty score means and standard errors for the 484 verbs of interest

To confirm this variation statistically, we ran a one-way analysis of variance (ANOVA). Using the R lme4 package (Bates et al. 2015), we fitted a linear mixed effects model of penalty scores, with verb as the predictor and random intercepts for participant (but not random slopes, due to model convergence problems, we suspect driven by the sheer size of the data sets). We then ran an ANOVA on this model, calculating p -values with the Satterthwaite approximation in the lmerTest package (Kuznetsova et al. 2017). We found a significant effect of verb on penalty scores, without and with prior context (without prior context, $F(483, 13658)=4.41$; $p<.001$; with prior context, $F(483, 6414)=3.24$; $p<.001$). Put differently, penalty scores show significant between-verb differences, as expected if bridge effects exist both without and with context.

5.2. Do bridge effects persist after controlling for a specific predictor?

We ran a second analysis closely modeled upon Liu et al.’s analyses. In theory, this second analysis addresses the same question – do penalties differ by verb? But this analysis differs from the analysis in Section 5.1 in two ways: (i) it uses individual judgments of each condition as an outcome variable rather than penalty scores (which likely leaves more variance for the model to partition), and (ii) it uses a theoretical predictor (backgroundedness, frame frequency, or semantic similarity) instead of the atheoretical predictor “verb” (which allows us to see if the short and long conditions behave differently in the presence of the predictor). We fitted linear mixed effects models predicting sentence acceptability based on crossing two factors: wh-dependency length (short, long) and the continuous independent measure of interest for each of the three prominent theories. For thoroughness, we constructed a different model for each frequency measure as well as each information structure and semantic similarity measure. To the extent that we find interaction effects, that would imply that the acceptability of sentences with clause-embedding verbs cannot be simply modeled using just two main effects (namely, wh-dependency length and each of the various measures), contra Liu et al.

Each model included only by-participant and by-item random intercepts, as not all models converged with random slopes. We take these models to provide evidence for bridge effects when

the interaction term’s p -value is less than .05 (according to lmerTest); and the model’s Bayes factor relative to a model without an interaction term (BF_{10}) is above 3 (calculated with the bayestestR package; Makowski et al. 2019), i.e. the data is at least 3 times more likely under a model with the interaction term than without an interaction term (Jeffreys 1961, Kass & Raftery 1995).

As Table 1 shows, significant interaction effects were reliably detected in our data, indicating bridge effects. Figure 2 illustrates the interactions, by plotting one selected measure for each theory. For information structure and frequency-based theories, both “no context” and “with context” data sets have interactions (in the predicted direction) with p -values below .05 and Bayes factors above 3. For template-based theories, the two GloVe models and the Wordnet model meet these criteria; only the LSA-based models do not. Overall, we take these results as converging evidence that there are bridge effects as classically defined in the literature (contra Liu et al. 2022).³

	No prior context					With prior context				
	b (s.e.)	t	p	Eff. Size	BF_{10}	b (s.e.)	t	p	Eff. Size	BF_{10}
Information structure theory										
% True responses	-0.33 (0.02)	-14.02	<.01	0.33	>100	-0.39 (0.03)	-13.72	<.01	0.39	>100
% Not-False responses	-0.74 (0.05)	-13.51	<.01	0.51	>100	-0.77 (0.07)	-11.28	<.01	0.53	>100
Template-based processing theory (hybrid score)										
LSA/Wikipedia	-0.12 (0.04)	-3.01	<.01	0.11	0.5	-0.09 (0.05)	-1.68	.09	0.07	<0.1
GloVe/Wikipedia	0.23 (0.03)	6.50	<.01	0.26	>100	0.35 (0.04)	8.21	<.01	0.40	>100
GloVe/Gigaword	0.17 (0.04)	4.84	<.01	0.19	>100	0.30 (0.04)	6.69	<.01	0.33	>100
Wordnet (log)	0.15 (0.03)	4.52	<.01	0.17	>100	0.16 (0.04)	3.90	<.01	0.19	13.5
Frequency-based processing theory										
Frame frequency (log)	0.06 (0.01)	9.67	<.01	0.30	>100	0.08 (0.01)	11.48	<.01	0.45	>100
Verb bias (log)	0.12 (0.01)	11.37	<.01	0.37	>100	0.19 (0.01)	14.10	<.01	0.56	>100

Column labels: “b”, “s.e.”, “t”, “p” = coefficient, standard error, t -value, and p -value of the interaction term, according to lmerTest (with a floor of .01); “Eff(ect) size” = cumulative size of the interaction effect over the full range of each independent measure (in z-units); “ BF_{10} ” = probability of the data under the model with an interaction term relative to a model without this term (with a ceiling of 100 and a floor of 0.1). Analyses with untransformed values were also run but produced worse model fits; these are not reported here for space reasons. Similarly, for the template-based processing theory, the only results reported are for the hybrid similarity score; model fits for *say* and *think* similarity scores are roughly equal (see Appendix C).

Table 1: Estimates for interaction effect for models of acceptability, with or without prior context

³ Liu et al. (2022) created ordinal mixed effects models on acceptability ratings, whereas we ran linear mixed effects models on z-transformed Likert ratings. Out of an abundance of caution, we also fitted three ordinal mixed effects models of “no context” raw ratings for the same set of verbs, with wh-dependency length as one fixed effect and a representative predictor for each of the theories – %True responses, GloVe/Wikipedia semantic similarity, and log frame frequency – as the other fixed effect. All three models detected significant interaction effects in the predicted direction, and the information structure and frequency models had Bayes factors above 3.

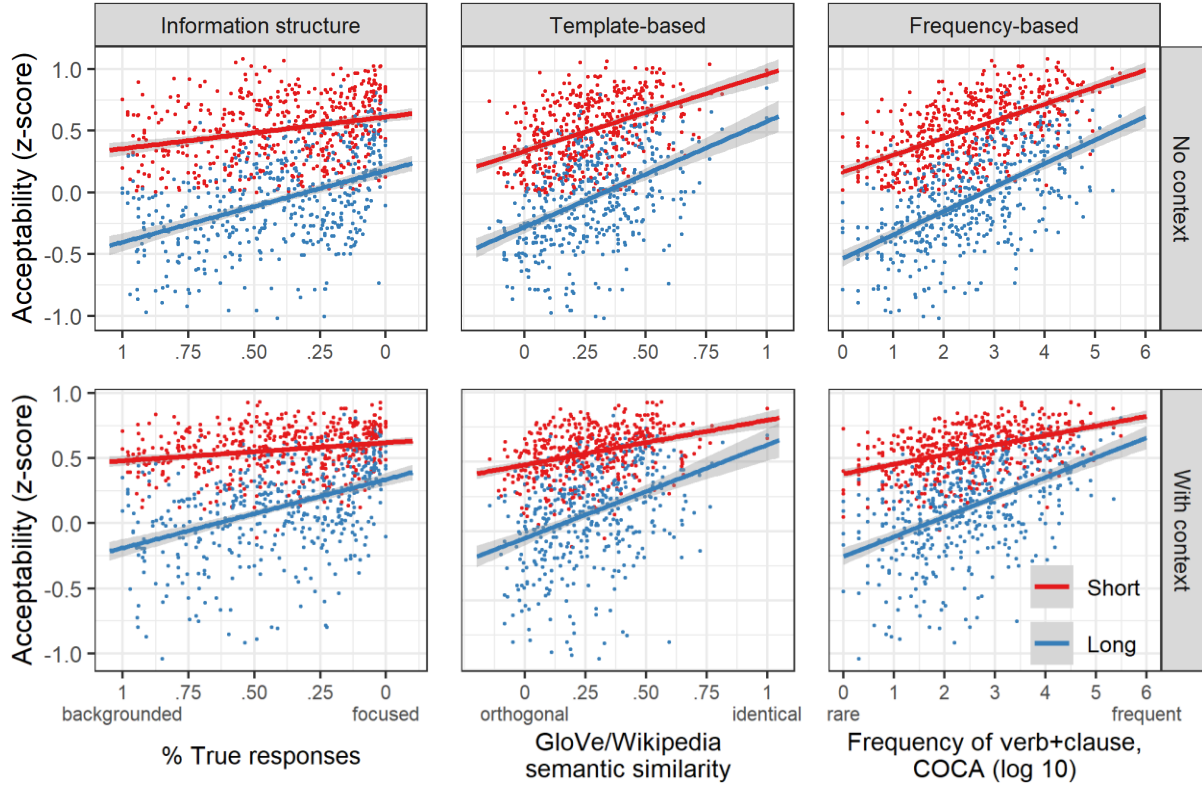


Figure 2: Interaction plots for selected predictors of information structure, template-based processing, and frequency (each dot = one verb)

We suspect Liu et al.’s null result (compared to our positive result) reflects a sample size difference: 48 verbs vs 484 verbs. To test this, we ran resampling simulations. We created 5,000 random samples of 48, 100, 200, 300, and 400 verbs, and fitted interaction models of z-scored “no context” acceptability ratings for each set, crossing dependency length and log frame frequency. Each model contained by-participant and by-verb random intercepts (so that almost all models would converge). Figure 3 is a box-and-whisker plot of the t -values of the interaction term for each sample size (showing the median, 1st, and 3rd quartiles). What we see is that smaller sample sizes lead to smaller t -values: for a sample size of 48, 30% of the simulations had t -values below 2, which is often taken as a minimum threshold for statistical significance (e.g. Baayen et al. 2008).




Figure 3: *t*-values of interaction effect of frame frequency and dependency length for random samples of verbs (convergent models only)

5.3. How does context affect bridge effects?

The finding that bridge effects (i.e., variation in penalties between verbs) exist for both the “no context” and “with context” data sets suggests that context, at least as operationalized in a dialogue format, does not eliminate bridge effects, contrary to what one might expect given reports in information structure accounts about context improving wh-extraction. That said, there is a small increase in the acceptability of long wh-dependencies and correspondingly a small decrease in the mean size of penalties between “no context” and “with context” (long wh-dependencies: -0.05 z-units vs. 0.13 z-units; penalties: 0.56 z-units vs. 0.44 z-units). This can be seen as consistent with previous reports that context improves extraction from the complement clauses of non-bridge verbs (e.g. Erteschik-Shir 1973; Ambridge & Goldberg 2008; Chaves & Putnam 2020), although the magnitudes suggest that the amelioration is much more modest for a full set of verbs.

Figure 4 corroborates these conclusions, by plotting penalty scores for each of the 484 verbs of interest, along with a line of best fit (solid) for the correlation between the two data sets, and a dashed line showing a hypothetical slope of 1.




Figure 4: “With context” and “no context” penalties (one dot = one verb)

We find a strong correlation ($r(482) = .623$, $p < .01$); if context eliminated bridge effects, the “with context” penalties would be centered around the same mean (plus random variation) and therefore the correlation should be closer to 0. The fact that the slope of the line of best fit is below 1 indicates that “with context” penalties tend to be smaller than “no context” penalties.

Taken together, these observations suggest that context reduces penalties to some degree, but this reduction applies either across-the-board such that the overall variation in penalties between verbs remains unchanged, or disproportionately on certain verbs so that overall variation increases.

5.4. The size of bridge effects

Finally, though our results confirm that there are in fact bridge effects in English, and that they do not decrease with context, we note that the mean size of bridge effects is relatively small for all of the measures tested here: about 0.2 to 0.5 z-scores over the entire range of the measures for “no context” penalties (Table 1). These effect sizes are smaller than typical island effects in English, which range from 0.6 to 1.2 z-scores in a recent review (Sprouse & Villata 2021). Given

that bridge effects and island effects are sometimes analyzed as related locality phenomena (e.g. Erteschik-Shir 1973, Ambridge & Goldberg 2008, also Goldberg 2006, Chaves & Putnam 2020, and references cited therein), this effect size difference may be meaningful. We explore this possibility in Section 8.1.

6. Evaluating the source of bridge effects

Having established that bridge effects exist both without and with context, we can move to the central theoretical debate in the literature: What is the source of bridge effects? In order to do so, we constructed a set of linear regression models for each existing theory of bridge effects based on the approach in Ambridge & Goldberg 2008: the dependent variable is the mean penalty for each verb and the sole predictor is a measure of interest from each theory. As above, we repeat the analyses for “no context” and “with context” penalties.

We will report following aspects of each regression model:

1. The **regression slope**, which indicates the size and direction of the effect of the measure of interest. The regression slope should closely track the interaction coefficients reported in Section 5.
2. The **Bayes factor of each regression model** (BF_{10}). BF_{10} indicates the ratio of the likelihood of the data under the experimental hypothesis to the likelihood of the data under the null hypothesis. BF s greater than 3 are conventionally interpreted as meaningful evidence for the experimental hypothesis.
3. **R^2** , which indicates the proportion of the variance in penalties that is explained by the measure of interest (ranging from 0 to 1).
4. **R^2 , corrected for attenuation**. The strength of the relationship between two variables is constrained by the reliability of each variable (Spearman 1904, Muchinsky 1996, etc.). This means that the lower the reliability, the lower the R^2 s. We report both raw R^2 s and R^2 s corrected for this attenuation (using Spearman’s method). We estimated reliability for acceptability judgments and backgroundedness measures using a bootstrap-based resampling simulation (see Appendix B for details). Because this resampling method is not feasible for the frequency and semantic similarity measures, we assume perfect reliability for them.

6.1. Evaluating the three theories

Figure 5 plots the relationship between penalties and one representative predictor per theory, and Table 2 reports the regression results for the full set of predictors.

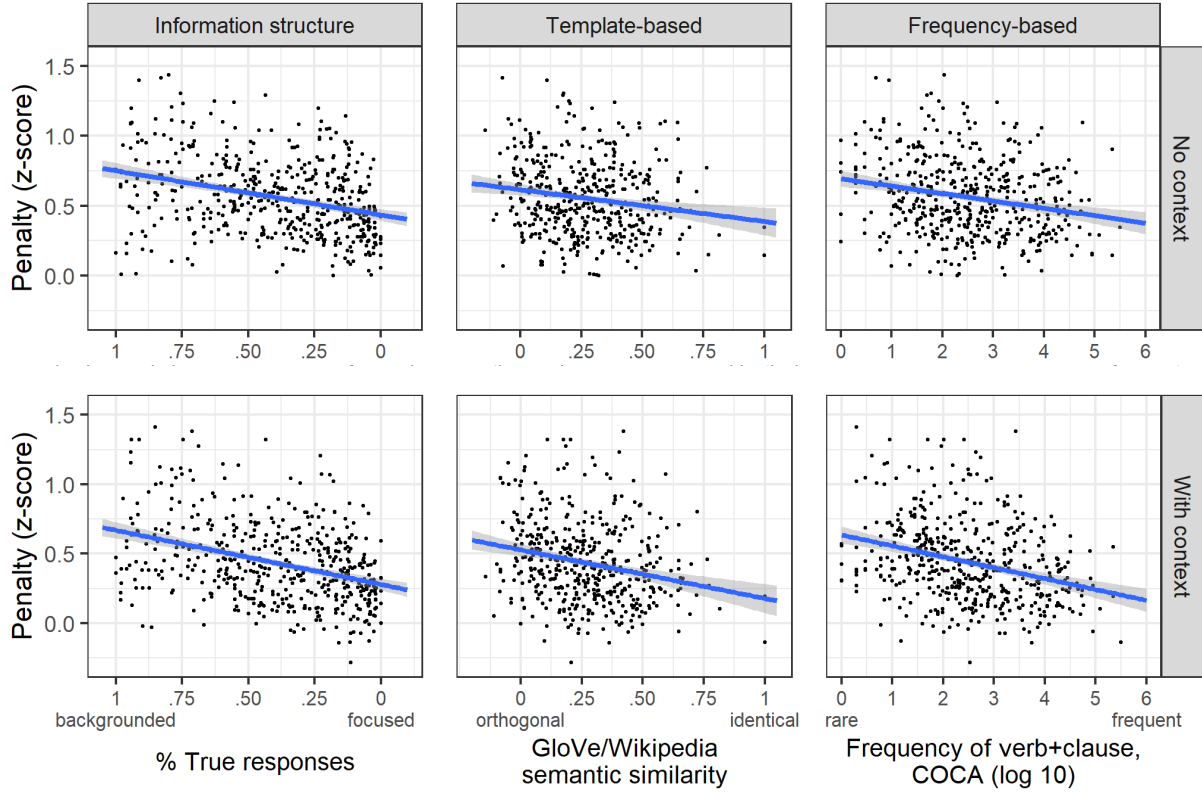


Figure 5: Long-distance penalty plots for selected predictors of information structure, template-based processing, and frequency.

	No prior context						With prior context					
	<i>b</i> (s.e.)	<i>t</i>	Eff. Size	BF ₁₀	Corr. R ²	Corr. R ²	<i>b</i> (s.e.)	<i>t</i>	Eff. Size	BF ₁₀	Corr. R ²	Corr. R ²
Information structure theory												
% True	0.32 (0.04)	7.41	0.32	>100	.103	.133	0.39 (0.05)	8.62	0.39	>100	.134	.186
% Not-False	0.74 (0.10)	7.22	0.51	>100	.098	.143	0.75 (0.11)	6.95	0.52	>100	.091	.143
Template-based processing theory (hybrid score)												
LSA/Wikipedia	0.11 (0.08)	1.44	0.10	0.1	.005	.006	0.08 (0.09)	0.95	0.07	0.1	.002	.003
GloVe/Wikipedia	-0.23 (0.07)	-3.52	0.27	21.8	.027	.034	-0.35 (0.07)	-5.03	0.40	>100	.054	.072
GloVe/Gigaword	-0.18 (0.07)	-2.65	0.20	1.6	.015	.019	-0.30 (0.07)	-4.06	0.33	>100	.036	.047
WordNet path-similarity (log)	-0.15 (0.07)	-2.32	0.17	0.7	.011	.014	-0.17 (0.07)	-2.38	0.19	0.8	.012	.016
Frequency-based processing theory												
Frame frequency (log)	-0.05 (0.01)	-4.92	0.30	>100	.048	.060	-0.08 (0.01)	-6.92	0.43	>100	.091	.121
Verb bias (log)	-0.12 (0.02)	-5.78	0.35	>100	.065	.081	-0.18 (0.02)	-8.70	0.55	>100	.137	.181

Table 2: Comparison of models of various theories of the source of bridge effects

We first discuss the “no context” data set, as previous experiments on bridge effects have focused on wh-questions presented without context.

Effect sizes are comparable to the results in Section 5 (as expected). Bayes Factors are high for all information structure and frequency predictors and for only one of the template-based processing predictors (GloVe/Wikipedia).

Turning to R^2 s, we see that the template-based processing theory explains very little variance, about 0.5% to 2.7% before correction, and 0.6% to 3.4% after correction. This accords well with our Bayes Factors results and with Ambridge & Goldberg's (2008) findings for a sample of 12. Here we expand Ambridge & Goldberg's result to four similarity measures and 484 verbs, and confirm that semantic similarity is a relatively poor predictor of bridge effects.

The frequency-based processing theory fares only slightly better. Depending on the measure, frequency accounts for 4.8% or 6.5% of variance before correction, and 6.0% or 8.1% after correction.

The information structure theory explains the most variance – 9.8% or 10.3%, depending on the measure, before correction, and 13.3% or 14.3% after correction. However, these values are substantially lower than the uncorrected R^2 of .69 that Ambridge & Goldberg (2008) observed for a sample of 12 verbs (note that this difference would likely be even more extreme if the Ambridge & Goldberg R^2 were corrected, which would likely increase the R^2). Though the interpretation of R^2 is subjective, given that Ambridge & Goldberg argued in favor of the information structure theory on the basis of explaining 69% of the variance, we suspect that 13.3% to 14.3% would be seen as relatively underwhelming evidence for the information structure theory.

We believe that Ambridge & Goldberg's substantially larger R^2 likely reflects the well-known fact that smaller samples with power issues – a possibility that they acknowledge of their study – can overestimate the effect size and hence r/R^2 . To test this, we replicated their backgroundedness analysis by analyzing the same 12 verbs from our “no context” data set. The uncorrected R^2 is .485, substantially higher than the R^2 for our full data set and much closer to Ambridge & Goldberg's R^2 .

We turn next to the “with context” results. Effect sizes and Bayes Factors are comparable to the “no context” results. Uncorrected R^2 s show a small increase, resulting in a similar range of variance accounted for: 0.2% to 13.7%. We see a larger increase in the corrected R^2 (because of slightly lower reliability in the smaller “with context” data set). The template-based processing theory still performs relatively poorly, accounting for 0.3% to 7.2% of the variance (post-correction). Both the frequency-based and information structure theories perform better: 12.1% and 18.1% of the variance for frequency, and 18.6% and 14.3% for information structure. This remains markedly lower than the Ambridge & Goldberg benchmark, and to our minds, lower than the field would expect for a good theory of bridge effects.

6.2. Are R^2 s low because of noise from infrequent verbs?

In testing a nearly exhaustive set of verbs, we necessarily included many low-frequency verbs. One possibility is that these verbs were unfamiliar to participants, and so penalty scores for these verbs would be unreliable, which could distort model outcomes and reduce R^2 .

We think such a scenario is unlikely. First, as noted in Section 3, we only analyzed responses from participants who demonstrated familiarity with the verbs on the synonym post-test.

Second, we find no evidence that penalty scores for infrequent verbs are more variable (unreliable) than frequent verbs. To test this, we calculated the standard deviation of “no context” penalties for each verb, obtained the frequency of each verb in COCA (regardless of the presence of complement clauses), and calculated a correlation. We did not see a significant (negative) correlation between standard deviation and frequency ($r(482)=-.036, p=.43$).

Third, we explored how much R^2 would improve if we excluded infrequent verbs. We sorted the verbs into twenty bins based on their COCA frequency. For the best-performing predictors of each theory – %True responses (information structure), GloVe/Wikipedia similarity (template-based), and log verb bias (frequency-based) – we ran twenty simple linear regression models, incrementally leaving out the low-frequency bins each time, i.e. first fitting data for all verbs, and then for verbs above the 5th, 10th, 15th, etc., percentiles. Figure 6 plots how R^2 changes with this frequency threshold. These values are the points in the plot. There are marked improvements in R^2 but only when we analyze the most frequent 40% of verbs (or higher).

However, there is also a general tendency for R^2 ranges to widen (often producing larger estimates) as sample sizes shrink. To account for this, we ran a Monte Carlo simulation in which verbs were randomly sorted into 20 bins to mimic the percentile analysis above, repeated 5,000 times. We calculated 95% intervals of R^2 based on this simulation, and plotted them as the gray bands in Figure 6. As this figure illustrates, the R^2 increase observed when we restrict the analysis to more frequent verbs almost always lies within these 95% intervals. This suggests that the increase is not uniquely attributable to verb frequency, but also attributable to sample size.




Figure 6: R^2 s for subsets of clause-embedding verbs above selected frequency percentiles, for best-performing predictors of “no context” bridge effects, with 95% intervals based on the random assignment of verbs into 20 bins.

These analyses suggest that the underperformance of existing theories cannot be meaningfully attributed to the many low-frequency verbs in our data set.

6.3. Takeaways for the debate on the source of bridge effects

In this section, we leveraged our new data sets to attempt to resolve the central debate in this literature: what is the source of bridge effects? We found that the information structure-based theory of bridge effects performs slightly better than the frequency-based and template-based processing theories. However, we also identified two challenges. First, none of these theories provide particularly strong fits to the full set of verbs. Second, adding context had little impact on effect sizes and slightly increased model fits for existing theories. This second outcome presents complications for all three theories. Frequency-based and template-based processing theories are typically silent about the role of context, implying that prior context should cause no change in bridge effects, and therefore no change in the correlation with predictor measures or R^2 s. Information structure theories, on the other hand, report that context can decrease penalties for at least some verbs, inviting the inference that bridge effects might decrease or even disappear given a supportive context.

These findings suggest that research on bridge effects could benefit by considering a wider range of theories beyond these three single-predictor theories. In the next section, we attempt to do just that.

7. A theoretical path forward: Morphosyntactic licensing combined with semantic/pragmatic constraints and processing costs

Our initial goal for this study was confirmatory: to test the empirical predictions of existing theories of bridge effects on a nearly exhaustive set of verbs. However, our results suggest that none of the theories perform particularly well. Therefore, in this section we add an exploratory goal: to find a new theoretical path forward based on the evidence made available in the new data sets. We do this in two steps. The first is to leverage our data sets to identify a new theoretically-relevant predictor for bridge effects. The second step is to integrate this new predictor with existing predictors in a theoretically-consistent way. More specifically, we draw upon suggestions from leading work on bridge effects (in particular Erteschik-Shir 1973, Richter & Chaves 2020, and Chaves & Putnam 2020), and propose a multivariate, layered account in which certain clause-embedding verbs can license long-distance wh-extraction, via subcategorization; in addition, the information structure, semantic, and frequency properties of verbs can further influence the acceptability of a wh-dependency. We show that this layered view of bridge effects delivers a substantially improved fit of our data compared to the various single-predictor theories that we have been considering so far, even after accounting for complexity. This finding provides new empirical evidence in favor of a multivariate approach toward bridge effects and also suggests that syntax has an important role to play in this multivariate approach.

7.1. Morphosyntactic licensing

Our new data sets [will be] available for all researchers to search for new potential predictors for bridge effects. In our own inspection of penalty scores, we noticed that verbs that have low penalty scores tend to be verbs that allow both a finite clausal complement and some form of nonfinite complement, such as exceptional case marking (e.g. *believe*) or control verbs (e.g. *claim*). To our knowledge, this has not previously been noted. To explore this further, we annotated our full list of verbs with subcategorization information (based on Levin 1993 and our own judgments). Figure 7 shows the distribution of penalties based on subcategorization for nonfinite complements.




Figure 7: Long-distance penalties of verbs allowing/not allowing nonfinite complement clauses

We first calculated a simple linear regression to predict penalties based on nonfinite complementation. Though nonfinite complementation is a categorical predictor, the model fit as indicated by R^2 is higher (.159) than the existing theories for the “no context” penalty scores and relatively similar (.127) to the best-performing existing theories for the “with context” penalty scores. Figure 8 illustrates this by plotting the uncorrected R^2 values for the best-performing predictors of the existing theories and nonfinite complementation.




Figure 8. *Uncorrected model fits for best-performing predictors of current theories and nonfinite complementation*

We take this as a potentially interesting research direction. As reviewed in Section 2, existing theories of bridge effects have mostly focused on non-syntactic factors under the assumption that the syntactic structure of all finite-clause embedding verbs is identical (with the exception of manner-of-speaking and/or factive verbs, e.g., Stowell 1981, Snyder 1992, Kastner 2015, Stoica 2016, de Cuba 2018, among others; and also see Kiparsky & Kiparsky 1970 for a complex NP analysis of factive verbs, although they explicitly reject this analysis). Our analysis here suggests it may be profitable to explore a syntactic difference across a wider range of verbs.

The question, of course, is how subcategorization for non-finite complements could be relevant for extraction from finite complements. One possibility is to connect it to recent cross-linguistic research on A-dependencies that cross finite clause boundaries, like exceptional case marking (ECM) and indexical shift (e.g., Wurmbrand 2019, also 2018). Briefly, in these dependencies, the subject of a verb’s (finite) complement clause behaves as if it were syntactically related to the main clause. Based on a cross-linguistic survey, Wurmbrand suggests that this is because the subject comes to occupy a special position in the complement clause’s left periphery, which is high enough for the subject to enter into dependencies with elements in the main clause. In ECM, for instance, the subject is structurally high enough to get case-marked by the main clause’s verb.

While Wurmbrand (2019) does not give a specific analysis of English ECM or other constructions with nonfinite complement clauses (like control), extending her proposal to these constructions seems not unfeasible. One possibility would be for English verbs that subcategorize for both finite and non-finite complements to have a similar special position present in both complements, and for wh-phrases to exploit this position to license further extraction to the left edge of the matrix clause (cf. the idea of “escape hatches” in Chomsky 1977, among many others). In other words, the morphosyntax of certain verbs licenses the possibility of extraction (out of all clausal complements for that verb). There are undoubtedly additional predictions made by this

analysis that could be explored in future work. But for this section, we will focus on embedding this morphosyntactic licensing property within a theory that can better predict bridge effects.

7.2. Adding semantic/pragmatic constraints and processing costs

Nonfinite complementation alone provides at best a slightly improved fit for bridge effects over the other single-source theories, and so it is unlikely to constitute a complete theory of bridge effects by itself. But it could be part of a layered theory in which extraction is licensed by morphosyntax (e.g., the availability of escape hatches), but is additionally constrained by the information structure constraints and processing costs proposed in existing theories. We take it to be uncontroversial that theories that propose a syntactic component to licensing long-distance dependencies also assume that other factors, such as information structure and processing complexity, continue to impact acceptability (though, the reverse may not be true – some theories of information structure or processing complexity may eliminate the need for a syntactic component to explain acceptability). Furthermore, we see this approach as drawing on several prior suggestions from the bridge effects literature. First, in order to explain cross-linguistic variation in extraction, Erteschik-Shir (1973) proposes that there is a class of “potential bridges” based on information structure, and that a subset of those become actual bridges through the acquisition process. A theory that combines morphosyntactic licensing and other constraints shares the same hierarchical arrangement (albeit working with different theoretical primitives), and provides similar flexibility to capture cross-linguistic variation (though, perhaps moving it to the syntactic component). Second, after concluding that the frequency-based theory is inadequate, Richter & Chaves (2020) suggest that combining semantic and pragmatic factors could provide a better explanation for bridge effects. Here we, too, are suggesting to combine multiple types of factors. Finally, Chaves & Putnam (2020) argue that locality in general (encompassing both bridge effects and island effect) might best be explained with an “eclectic” theory that draws on syntactic, semantic/pragmatic, and processing factors. What we do here is to develop a specific version of this kind of approach.

To empirically evaluate this layered theory, we created six regression models that instantiate different theoretically-driven possibilities, comparing both R^2 s (uncorrected for attenuation) and Bayesian Information Criterion (BIC), which balances data fit with a penalty for increased complexity (the lower the BIC, the better) (Figure 9). We include four single-predictor theories as baselines: information structure, template-based, frequency, and syntax (nonfinite complementation). To facilitate comparison, we selected the best-performing predictor for each non-syntactic theory and a set of 442 verbs that have complete predictor information for these theories. We then consider two theories that combine predictors: a model that combines only the three non-syntactic predictors, which we call *multivariate, no syntax*, and a model that combines all four predictors (including nonfinite complementation), which we call *multivariate, with syntax*. Both multivariate models are additive, meaning there are no interaction terms in the models. We believe this is consistent with existing theories: each component is independent of the others.⁴

⁴ Out of an abundance of caution, we also ran the full interaction models. Interaction models always have higher BICs than the additive models, suggesting that the increased fit of the interaction models (between .01 to .05 in R^2) is outweighed by their greater complexity.




Figure 9. Model fits for multivariate models of bridge effects

For both “no context” and “with context” penalties, the “multivariate, with syntax” model has the highest R^2 s of all the models under consideration, including the “multivariate, no syntax” model. Crucially, the R^2 s are substantially higher, in the .25-.3 range, or about 2.5 times the R^2 s for the best-performing single-predictor theory in the literature (namely, information structure). Furthermore, even though the “multivariate, with syntax” model is the most complex of the various models, it has the lowest BIC values, confirming that it achieves the best coverage of the data, even after controlling for complexity.

There are potentially interesting orderings among the other theories, such as the fact that the syntax-only model outperforms the “multivariate, no syntax” model for “no context”, and that R^2 s for the “multivariate, no syntax” model increases substantially for “with context” (a point we consider in section 8.2). But we take the primary result of this analysis to be that a layered theory that features syntactic licensing in addition to information structure constraints and processing costs, as proposed here, yields substantially better model fits than single-predictor theories or a multivariate non-syntactic theory, even after taking into consideration its relative complexity. That said, it is an open question whether our results satisfy the field’s conception of a good theory of bridge effects (cf. Ambridge & Goldberg’s R^2 of .69, albeit for an underpowered 12-verb sample). We note that our R^2 values are an underestimate; actual R^2 s will almost certainly be higher after correcting for attenuation (but we do not do that here, because to our knowledge, there is no correction formula applicable for models with multiple predictors). It may also well be the case that there are other components in a layered theory that we have not yet uncovered. But we hope that this illustrates a promising path forward based on the new information in our data sets.

8. Implications beyond bridge effects

8.1. The unification of bridge and island effects in the theory of locality

Given that bridge effects are a type of locality phenomenon, one question that arises in the literature is whether bridge effects and islands effects can be unified under a single analysis, as

often proposed in information structure theories (e.g. Erteschik-Shir 1973, Ambridge & Goldberg 2008, and references cited therein; cf. Chaves & Putnam 2020). However, our results raise a novel empirical challenge for the unification of bridge effects and island effects: as Sections 5 and 6 noted, the effect sizes of bridge effects are substantially smaller than the effect sizes of island effects. For example, the maximum impact of backgroundedness on penalties – going from not backgrounded at all to totally backgrounded – is about 0.2-0.5 z-units depending on the measure and the data set. In contrast, various island effects in English that have been tested using similar experimental methods appear to have effect sizes of 0.6-1.2 z-units (see Sprouse & Villata 2021 for a review). This effect size difference suggests either that bridge effects should be treated as distinct from island effects, or that theories seeking to unify them, whether rooted in information structure or otherwise, need to include an additional layer of complexity to explain the differing effect sizes.

We are personally inclined to interpret the difference in effect sizes between bridge and island effects as reflecting different sources. That said, there is a three-way distinction in Erteschik-Shir’s (1973) theory that could potentially be adopted for a unified non-syntactic account in order to explain the effect sizes we observe. As alluded to in section 7.2, to explain cross-linguistic differences between Danish and English, Erteschik-Shir (*ibid.*:125ff) suggests that there is a set of “potential bridges”, which are presumably universal, while usage within a language determines which potential bridges become actual bridges. We note that this analysis implies at least three different classes of constituents – non-bridges, potential bridges that do not become actual bridges, and potential bridges that do. We note that this three-way distinction could be exploited to explain the observed effect sizes (although not necessarily consistent with Erteschik-Shir’s information structure proposal): island effects might correspond to extraction from non-bridges, while the variation found in bridge effects might correspond to extraction from potential bridges that do not become actual bridges (higher penalties) as well as extraction from actual bridges (lower penalties). Researchers interested in this approach could use our backgroundedness and acceptability data sets, perhaps by combining them with island effects data sets, or by collecting similar data sets in other languages to quantify the cross-linguistic variation in bridge effects.

8.2. The effect of context

We originally collected “with context” penalties in response to reports in information structure theories that supportive context can make long-distance wh-extraction more acceptable, which raise the possibility that context can decrease or eliminate bridge effects. Though we found no such impact on bridge effects in Section 5, we did observe a novel effect of context in the analyses in Sections 6 and 7: model fits (R^2 s) are higher for “with context” penalties for information structure, template-based, frequency-based, and multivariate theories (but not for nonfinite complementation when it is the sole predictor). These higher R^2 s warrant further exploration, both for what they could mean for theories of bridge effects, and for what they could mean for a more general theory of the effect of context on the acceptability of long-distance dependencies.

The first question we can ask is whether the pattern that we observed – the higher R^2 s – is meaningful – that is, whether it is beyond what we would expect due to sampling error between the two data sets. Sampling error is a particularly plausible explanation because our “with context” data set is about half the size of the “no context” data set in terms of the number of observations per verb (due to the addition of catch trials to ensure the context was read). To rule out this possibility, we ran a bootstrap analysis to determine if the “with context” R^2 s are more extreme than we would expect based on the “no context” data set. We randomly sampled participant

responses with replacement from the “no context” data set such that the size of the random sample for each verb is equal to the actual set of “with context” observations for that verb. We then calculated penalty scores from these random samples and fitted the same regression models as in Section 7.2. We repeated the process 5,000 times to generate expected distributions of R^2 values under repeated sampling. We then constructed a confidence interval with the 2.5th and 97.5th percentiles of each distribution. Observed “with context” R^2 s are higher than the upper bound of this interval for all but the syntax-only model (Figure 10), indicating that higher R^2 s for “with context” data are unlikely a sampling artifact.




Figure 10: Distribution of simulated R^2 s (excluding values beyond 2.5th and 97.5th percentiles) and observed “with context” R^2 s (diamonds)

Given that the increase in R^2 appears to be meaningful, we can next ask why context improves the fit of several (non-syntactic) predictors. One possibility is that there are actually two effects combined in the “no context” penalties: a non-syntactic (information structure, semantic or frequency) effect, as hypothesized in existing theories, and some additional pragmatic effect. Adding supportive context eliminates or reduces this second effect, providing a better estimate of bridge effects, and thus allowing the non-syntactic factors to perform better. Identifying this second effect is beyond the scope of this paper, as it will require a general theory of how context affects acceptability judgments. Though no such theory currently exists, the suggestions in information structure-based theories about the role of context may be a useful place to start (see works cited above). Furthermore, the data sets and predictors that we compiled here could be used to test theories about which semantic/pragmatic properties (of verbs or even the specific sentence frames that we constructed) are more likely to be affected by (dialogue) context.

9. Conclusion

Bridge effects have been variously attributed to information structure constraints or processing factors related to template-based processing or frequency effects (e.g. Erteschik-Shir 1973, Ambridge & Goldberg 2008, Dąbrowska 2008, Richter & Chaves 2020, Kothari 2008, Liu et al. 2022, among others). A recent study (Liu et al. 2022), pursuing an extreme version of a frequency-based processing account, has even suggested that bridge effects do not exist, once the frequency of a verb co-occurring with a finite complement clause is taken into account. The lack of a consensus on such basic questions around bridge effects, as we argued and as suggested in

recent work, partly reflects the fact that experimental studies on which these claims are based have studied relatively few clause-embedding verbs, which makes sampling errors more likely.

Our solution to this empirical and theoretical stalemate was to create two large-scale data sets of English bridge effects as benchmark data sets, which we [will make] publicly available. We collected acceptability judgments for sentences presented without and with prior context for a nearly exhaustive set of 640 clause-embedding verbs, and compiled theoretically-relevant measures for each verb, such as backgroundedness judgments, semantic similarity measures, and frequency estimates. Focusing on a subset of 484 verbs for whom finite complement clauses are most likely grammatical for our participants, we addressed three questions about bridge effects: Do they exist at all? Which existing theory best explains the source of bridge effects? And are there new patterns in our data sets that could lead to a better theory?

Across the full range of verbs, we found clear evidence of bridge effects: verbs do vary significantly in whether they allow long-distance wh-extraction. Bridge effects exist even after accounting for frequency (contra Liu et al. 2022), backgroundedness, and semantic similarity. Bridge effects were also observed in the presence of a dialogue that provided prior context, contrary to what one might expect given prior reports that context can make long-distance wh-extraction more acceptable for certain verbs.

With this confirmation of the existence of bridge effects, we then statistically evaluated the three leading theories about their source. We found that the information structure theory performs the best, but model fits for all three theories are relatively low, contrary to expectations and previously-reported experiment results (although these experiments had much smaller samples).

The underperformance of existing theories suggests that bridge effects could benefit from fresh theorizing. We identified from our data sets a novel morphosyntactic predictor of bridge effects – nonfinite complementation, which we believe potentially connects to a growing literature on cross-clausal A-dependencies in theoretical syntax (Wurmbrand 2019 and references therein). Integrating this new predictor with existing ideas in the bridge effects literature, we presented a multivariate layered theory of bridge effects in which wh-extraction is licensed by (morpho)syntax and is further subject to information structure constraints and processing costs. Our evaluation of this multivariate theory shows that it explains the largest share of observed variation (R^2), relative to a multivariate theory without a syntactic licensing component as well as single-predictor theories, in which bridge effects are attributable to a single factor (e.g. information structure, some processing factor, morphosyntax). The multivariate theory with a syntactic licensing component also has the lowest BIC, indicating that it achieves the best fit of the data (among the various models evaluated), after controlling for model complexity.

In sum, our investigation of a comprehensive set of English clause-embedding verbs provided novel evidence for bridge effects as a phenomenon to be studied, but showed that existing theories, elegant as they may be, provide only a limited explanation of overall variation. On a more positive note, our study here has pointed to a clear path forward. We demonstrated how insights from existing theories and a novel morphosyntactic predictor identified from our data sets can be combined fruitfully to develop a layered theory that offers better empirical coverage. In the longer term, we hope our data sets and findings will support efforts to develop more robust theories of bridge effects and wh-dependencies.

References

- Ambridge, Ben & Adele Goldberg. 2008. The island status of clausal complements: Evidence in favor of an information structure explanation. *Cognitive Linguistics* 19, 357–389.
- Ambridge, Ben, Julian M. Pine, and Elena V. M. Lieven. 2015. Explanatory adequacy is not enough: Response to commentators on ‘Child language acquisition: Why universal grammar doesn’t help’. *Language* 91(3): e116-e126.
- Anand, Pranav, Jane Grimshaw & Valentine Hacquard. 2019. Sentence embedding predicates, factivity and subjects. *Tokens of Meaning: Papers in Honor of Lauri Karttunen*, ed. by Cleo Condoravdi & Tracy Holloway King. Stanford, CA: CSLI Publications.
- Baayen, R. Harald, Douglas J. Davidson & Douglas M. Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59(4), 390–412.
- Bates, Douglas, Martin Mächler, Ben Bolker & Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67(1), 1–48.
- Biber, Douglas. 1999. A register perspective on grammar and discourse. *Discourse Studies* 1, 131–150.
- Bird, Steven, Ewan Klein & Edward Loper. 2009. *Natural language processing with Python: analyzing text with the Natural Language Toolkit*. O’Reilly Media.
- Chaves, Rui P. & Michael T. Putnam. 2020. *Unbounded dependency constructions: Theoretical and experimental perspectives*. Oxford: Oxford University Press.
- Chomsky, Noam. 1977. On wh-movement. In Peter Culicover, Thomas Wasow & Adrian Akmajian (eds.), *Formal syntax*, 71–132. New York: Academic Press.
- Fellbaum, Christiane. 1998. *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press.
- de Cuba, Carlos. 2018. Manner-of-speaking *that*-complements as close apposition structures. *Proceedings of the Linguistic Society of America* 3, 32-1.
- Dąbrowska, Ewa. 2008. Questions with long-distance dependencies: A usage-based perspective. *Cognitive Linguistics* 19, 391-425.
- Dąbrowska, Ewa. 2013. Functional constraints, usage, and mental grammars: A study of speakers’ intuitions about questions with long-distance dependencies. *Cognitive Linguistics* 24, 633-665.
- Davies, Mark. 2020. The Corpus of Contemporary American English (March 2020). Online: <https://www.english-corpora.org/coca/> (accessed 2019-2020).
- Dean, Janet. 1967. Noun phrase complementation in English and German. Ms., MIT.
- Deerwester, Scott, Susan T. Dumais, George W. Furnas, Thomas K. Landauer & Richard Harsham. 1990. Indexing by Latent Semantic Analysis. *Journal of the American Society for Information Science* 41, 391–407.
- Drummond, Alex. 2012. IbexFarm (version 0.3.7). [Software]
- Erteschik-Shir, Nomi. 1973. On the nature of island constraints. Ph.D. dissertation, MIT.
- Erteschik-Shir, Nomi. 2017. Bridge Phenomena. In Martin Everaert & Henk van Riemsdijk (eds.), *The Wiley Blackwell Companion to Syntax, 2nd Edition*. Oxford: John Wiley and Sons. <https://doi.org/10.1002/9781118358733.wbsyncom011>
- Fares, Murhaf, Andrey Kutuzov, Stephan Oepen & Erik Velldal. 2017. Word vectors, reuse, and replicability: Towards a community repository of large-text resources. *Proceedings of the 21st Nordic Conference on Computational Linguistics*, 271-276.
- Featherston, Samuel. 2004. Bridge verbs and V2 verbs -- the same thing in spades? *Zeitschrift für Sprachwissenschaft* 23, 181-209.

- Goldberg, Adele. 2006. *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press.
- Hale, John. 2001. A probabilistic Earley parser as a psycholinguistic model. *Proceedings of the second meeting of NAACL*, 159-166.
- Hofmeister, Paul & Ivan A. Sag. 2010. Cognitive constraints and island effects. *Language* 86(2), 366-415.
- Jeffreys, Harold. 1961. *Theory of probability* (3rd ed.). Oxford: Oxford University Press.
- Kass, Robert E. & Adrian E. Raftery. 1995. Bayes Factors. *Journal of the American Statistical Association* 90(430), 773-795.
- Kastner, Itamar. 2015. Factivity mirrors interpretation: The selectional requirements of presuppositional verbs. *Lingua* 164, 156-188.
- Kiparsky, Paul & Carol Kiparsky. 1970. Fact. In Manfred Bierwisch & Karl Erich Heidolph (eds.), *Progress in Linguistics*, 143-173. The Hague: Mouton.
- Kothari, Anubha. 2008. Frequency-based expectations and context influence bridge quality. In Michael Grosvald and Dionne Soares (eds.), *Proceedings of WECOL 2008: Western Conference on Linguistics*, 136-149. Fresno, CA: California State University.
- Kuznetsova, Alexandra, Per B. Brockhoff & Rune H. B. Christensen. 2017. lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software* 82, 1-26.
- Levin, Beth. 1993. *English verb classes and alternations: A preliminary investigation*. Chicago: University of Chicago Press.
- Levy, Roger. 2008. Expectation-based syntactic comprehension. *Cognition* 106, 1126-1177.
- Liu, Yingdong, Rachel Ryskin, Richard Futrell & Edward Gibson. 2022. A verb-frame frequency account of constraints on long-distance dependencies in English. *Cognition* 222, 104902.
- Makowski, Dominique, Mattan S. Ben-Shachar & Daniel Lüdtke. 2019. bayestestR: Describing effects and their uncertainty, existence and significance within the Bayesian framework. *Journal of Open Source Software* 4(4), 1541.
- Muchinsky, Paul M. 1996. The correction for attenuation. *Educational and Psychological Measurement* 56(1), 63-75.
- Müller, Sonja. 2015. Deriving island constraints with Searle and Grice. a pragmatic account of bridge verbs. *Studia Linguistica* 69(1), 1-57.
- Parker, Robert, David Graff, Junbo Kong, Ke Chen & Kazuaki Maeda. 2011. *English Gigaword, fifth edition* (LDC2011T07). Philadelphia, PA: Linguistic Data Consortium.
- Pennington, Jeffrey, Richard Socher & Christopher D. Manning. 2014. GloVe: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1532-1543.
- Richter, Stephanie & Rui P. Chaves. 2020. Investigating the role of verb frequency in factive and manner-of-speaking islands. *Proceedings of CogSci* 42, 1771-1777. Cognitive Science Society.
- Snyder, William. 1992. Wh-extraction and the lexical representation of verbs. Ms., MIT. Online: https://william-snyder.uconn.edu/wp-content/uploads/sites/2834/2019/11/Snyder_1992_Wh_V.pdf (accessed May 2023).
- Stowell, Timothy. 1981. Origins of phrase structure. Ph.D. dissertation, MIT.
- Spearman, Charles. 1904. The proof and measurement of association between two things. *American Journal of Psychology* 15, 72-101.
- Sprouse, Jon, Carson T. Schütze & Diogo Almeida. 2013. A comparison of informal and formal acceptability judgments using a random sample from *Linguistic Inquiry* 2001-2010. *Lingua* 134, 219-248.

- Sprouse, Jon, Beracah Yankama, Sagar Indurkha, Sandiway Fong & Robert C. Berwick. 2018. Colorless green ideas do sleep furiously: gradient acceptability and the nature of the grammar. *The Linguistic Review* 35(3), 575–599.
- Sprouse, Jon & Sandra Villata. 2021. Island effects. In Grant Goodall (ed.), *The Cambridge Handbook of Experimental Syntax*, 227–257. Cambridge: Cambridge University Press.
- Ștefănescu, Dan, Rajendra Banjade & Vasile Rus. 2014. Latent Semantic Analysis models on Wikipedia and TASA. *Proceedings of LREC 2014*, 1417-1422.
- Stoica, Irina. 2016. Island effects and complementizer omission: the view from manner of speaking verbs. In Petronia Petrar & Amelia Precup (eds.), *Constructions of Identity (VIII): Discourses in the English-Speaking World*, 191-200. Cluj-Napoca, România: Presa Universitară Clujeană.
- Uriagereka, Juan. 1999. Multiple spell-out. In Samuel D. Epstein & Norbert Hornstein (eds.), *Working minimalism*, 251-282. Cambridge, MA: MIT Press.
- Vasishth, Shravan, Daniela Mertzen, Lena A. Jäger & Andrew Gelman. 2018. The statistical significance filter leads to overoptimistic expectations of replicability. *Journal of Memory and Language* 103, 151-175.
- White, Aaron Steven & Kyle Rawlins. 2018. The role of veridicality and factivity in clause selection. In Sherry Hucklebridge & Max Nelson (eds.), *Proceedings of the 48th Annual Meeting of the North East Linguistic Society*, 221–234. Amherst, MA: GLSA Publications.
- White, Aaron Steven & Kyle Rawlins. 2020. Frequency, acceptability, and selection: a case study of clause-embedding. *Glossa* 5(1), 105.1-41.
- Wurmbrand, Susi. 2018. Cross-clausal A-dependencies. Lecture handouts and slides for course taught at the ACTL Summer School, London.
- Wurmbrand, Susi. 2019. Cross-clausal A-dependencies. In Eszter Ronai, Laura Stigliano & Yenan Sun (eds.), *Proceedings of the Fifty-Fourth Annual Meeting of the Chicago Linguistic Society*, 585-604. Chicago: Chicago Linguistic Society.