

Wh-the-hell as a polarity-insensitive, speaker-oriented domain restrictor*

Joshua Martin
Harvard University

Abstract I argue for a reanalysis of ‘wh-the-hell’ questions as restrictions on the domain of alternatives, rather than polarity items or domain extensions. This analysis returns to their original characterization as ‘aggressively non-D-linked’ by specifying what it means to be D-linked, and in doing so captures a number of otherwise unconnected properties of such questions without positing extraneous features. Consequences of the proposal include a new notion of question-specific contextuality, evidence for an inquisitive semantic condition on sluicing, and a potential general discourse restriction on wh-movement.

Keywords: ‘the hell’, D-linking, wh-questions, alternative semantics, polarity, context

1 Non-D-linked *wh*-questions

Imagine it is the day of an election, and the results are coming in live. It’s a close race, with potential legal challenges to the results. Your friend utters one of:¹

- (1) a. I’m freaking out about who won the election.
b. I’m freaking out about who the hell won the election.

The salient reading of (1a) is that they heard the final result, and are displeased with it, but there is also a reading on which it is the uncertainty of the result freaking them out. Interestingly, (1b) has only the latter reading, and is incompatible with a situation where the speaker knows the true victor. The same contrast holds in the unembedded context, where the non-*hell* question can be used to request currently unknown information or to express surprise at the answer (with echo intonation), while the *hell* question lacks the echo option and must be a genuine interrogative.

- (2) WHO (*the hell) won the election?

* I thank Kathryn Davidson, Isabelle Charneval, Gunnar Lund, Jonathan Bobaljik, C.T. James Huang, Uli Sauerland, Diti Bhadra, Hadas Kotek, and Harvard’s Meaning & Modality Lab for helpful discussions, as well as the reviewers and audience of SALT30 for their insightful questions.

¹ Unless otherwise noted, English judgments are from my own intuitions and checked informally with other native speakers. Any discrepancies in these judgments will be made explicit.

This paper aims to develop a semantics of *the hell* to account for its interpretive effects, like (1), and its syntactic restrictions, like (2). This class of *wh*-questions, exemplified by *wh-the-hell* but including also *what the fuck*, *what on earth*, and others, was first described by Pesetsky (1987) as ‘aggressively non-discourse-linked’. He argues that ‘the whole point of uttering a question like *What the hell did you read that in?* is to express surprise in the answer. The appropriate answer is presumed not to figure in previous discourse’ (Pesetsky 1987: 111). Not just indicating attitude, these constructions have a variety of distributional constraints not straightforwardly captured by such a discourse requirement. Along with being ungrammatical in echo questions as in (2), it is ungrammatical with *which* (Pesetsky 1987), embedded under some factive verbs (den Dikken & Giannakidou 2002; Huang & Ochi 2004), in-situ in single or multiple *wh*-questions (Obenauer 1994), when pied-piped (Pesetsky 1987; Akiyama 2014), and when sluiced (Merchant 2001: 122):

- (3) a. *Which the hell did you buy? e. *To who the hell did you sell
 b. *I know who the hell bought it. it?
 c. *Who bought what the hell? f. I wish I knew who the hell *(it
 d. *You bought what the hell? was)!

This collection of behaviors has led theorists to ascribe a wide variety of properties to *wh-the-hell*, including polarity sensitivity (den Dikken & Giannakidou 2002), identificational focus (Oguro 2017), and logophoric attitude (Huang & Ochi 2004; Akiyama 2014). While each of these accounts offers some useful insight, none of them has managed to capture the full empirical picture, I think *because* of their ‘add a new feature’ approach. In most of these accounts, ‘aggressively non-D-linked questions’ has become a category label, a way of referencing the class of *wh-the-hell* questions, while its contribution as an explanatory aspect has been systematically reduced. The question has become, roughly, ‘What other property do aggressively non-D-linked questions have that explains their felicity conditions?’. This approach seems misled. The question, as originally posed and as I raise it again, should rather be ‘What is it *about the aggressive non-D-linking of these questions* that explains their felicity conditions?’ A maximally parsimonious theory should explain the behavior of non-D-linked questions *in virtue of their being non-D-linked*.

My central claim will be that explaining *wh-the-hell* questions requires *only* explaining what it means to be non-D-linked. Since this requires formalizing a notion of D-linking, which has resisted formal specification despite its widespread use, *wh-the-hell* will act as a useful test case for illuminating our understanding of the syntax and semantics of *wh*-questions more broadly. In Section 2, I argue against the most widely cited account of *wh-the-hell*, that of den Dikken & Giannakidou (2002), and establish that *wh-the-hell* is not a polarity item. In Section 3, I develop

an analysis on which *the hell* restricts the domain of alternatives which can serve as possible answers to the question, and show that it captures most of the properties above, and some new ones. Section 4 addresses the persistent problem of in-situ *wh-the-hell* in multiple questions, proposing a general condition on *wh*-in-situ and linking it to the alternative semantics of questions, and Section 5 concludes.

2 Against polarity

A highly influential account of *wh-the-hell* from den Dikken & Giannakidou (2002), henceforth DG, posits that *wh-the-hell* is a polarity item (PI) based on observed parallels between it and known PIs like *any*. While it is theoretically attractive to assimilate *wh-the-hell* to the independently needed framework of polarity, it is ultimately untenable, as the parallels with *any* break down against a wider backdrop of data. DG's separate argument that *wh-the-hell* is a domain extender will be addressed in Section 3.5; here I focus primarily on the polarity arguments.

2.1 Veridicality and ignorance

Uncontroversially, *wh-the-hell* phrases are restricted with respect to which verbs they may embed under. DG observe that they are ungrammatical under factive or veridical verbs, but become acceptable when those verbs are negated:

- (4) a. I *(don't) know who the hell would buy that book. (DG: 33)
 b. He {didn't tell / *told} me who the hell bought that book.

This parallels the distribution of the PI *anyone*:

- (5) a. I *(don't) know that anyone would buy that book. (adapted from DG: 34)
 b. He {didn't tell / *told} me that anyone bought that book.

And, both *the hell* and the PI are licensed in many nonveridical environments:

- (6) a. I wonder who the hell bought that book. (adapted from DG: 34)
 ... if anyone ...
 b. He refused to tell me who the hell bought that book.
 ... if anyone ...
 c. Nobody knows who the hell bought that book.
 ... if anyone ...
 d. Only he knows who the hell bought that book.
 ... if anyone ...

- e. If he knows who the hell bought that book, he should tell us now.
... if anyone ...

DG conclude from this that their licensing conditions are the same, and thus that this data argues in favor of *wh-the-hell* being a PI. However, note that the only predicates which rule out *wh-the-hell* are the positive factive ones, *I know...* and *He told me...* All of the other environments share another discourse feature: they ensure that the speaker does not know the answer to the question. If I am wondering who bought the book, or you refuse to tell me, or nobody knows, or only you know, it is clear that I do not know. This factor is obviously irrelevant for licensing a PI, as we can illustrate by swapping the first and third-person pronouns.

- (7) a. He doesn't know that anyone would buy that book.
b. I didn't tell him that anyone bought that book.
- (8) a. He wonders if anyone bought that book.
b. I refused to tell him if anyone bought that book.
c. Nobody but me knows if anyone bought that book.
d. Only I know whether anyone bought that book.
e. If I know if anyone bought that book, I should tell him now.

Crucially, however, it is relevant for *wh-the-hell*. Not every context will show this, since it may be difficult to *guarantee* first-person ignorance of the answer, e.g. *He wonders...* says nothing about the speaker's knowledge either way. Still, when the environment implies the speaker does know the answer, *wh-the-hell* is ruled out.

- (9) a. *I refused to tell him who the hell bought that book.
b. *Nobody but me knows who the hell bought that book.
c. *Only I know who the hell bought that book.

Guaranteeing speaker ignorance, such as through contrastive focus as in (10), also allows *wh-the-hell* (but not *any*) to embed under a non-negated factive.

- (10) a. HE knows who the hell bought that book, but he won't tell me.
... *that anyone ...
b. HE told his brother who the hell bought that book, but not me.
... *that anyone ...

No theory of polarity predicts sensitivity to this. Per DG, (10) alone is enough to rule out *wh-the-hell* being a PI: 'there is one distributional constraint that all the PI paradigms must obey in order to qualify for PI status: they must be inadmissible in

the scope of veridical operators.’ (DG: 37) The only exception is another purported parallel between *wh-the-hell* and *any*, namely that they can be under factives when there is a ‘negative implicature’ (Giannakidou 1999):

- (11) NOW I know who the hell stole my car. (DG: 39)
Negative implicature: I did not know before.
- (12) I am surprised she has any friends. (DG: 38)
Negative implicature: I didn’t expect her to have any friends.

However, this conflates two flavors of ‘negative implicature’ which should be regarded as distinct. In particular, the ‘I did not know before’ implicature of focused *NOW* licenses *wh-the-hell*, but not *any*, as in (13). Conversely, the ‘I didn’t expect’ implicature of *surprise* licenses *any*, but not *wh-the-hell*, as in (14).

- (13) *NOW I know (that) she has any friends.
Negative implicature: I did not know before.
- (14) *I am surprised (by) who the hell stole my car.
Negative implicature: I didn’t expect them to have stolen my car.

It is thus incorrect to conclude ‘that *wh-the-hell* behaves like other PIs in this respect’ (DG: 39). This provides a double dissociation between *wh-the-hell* and *any*: we have shown environments where the former is licensed but not the latter, such as (10) and (11) vs. (13), and where the latter is licensed but not the former, such as (8) vs. (9) and (12) vs. (14). Regardless of whether analyzing *wh-the-hell* as a PI would capture its other behavior, then, the path is too hazardous to tread.

2.2 Negative attitude

Briefly, DG also give *wh-the-hell* a presupposition of negative attitude, ‘a modalized conditional statement saying that if any x such as x did what is expressed by the VP, then x should not have done it’ (DG: 43). So, the presupposition of (15a) is (15b).

- (15) a. Who the hell talked to Ariadne? (den Dikken & Giannakidou 2002: 43)
 b. If there is a person x in w , and x talked to Ariadne in w : x should not have talked to Ariadne in w .

But this is too strong. It is easy to imagine contexts where (15b) does not hold, e.g. if every student in a class was supposed to be interviewed and have a report written about them, but the report about Ariadne is missing, even though she insists someone did interview her, a teacher might utter (15a) with the intention that someone absolutely should have. In contexts like our election example (1), it is not

the case that you don't think anyone should win the election. Positive attitude is also possible; imagine that you've failed to get tickets to see a sold-out show but I miraculously found some at the last minute. You might say *How the hell did you do that?* and be quite happy that I did so. This is all just to say that negative attitude should not be encoded into *wh-the-hell*. It can be derived in many cases, simply from the interaction of wanting to know the answer and being unable to find it.

3 Semantic analysis

In this section, I build an analysis of *wh-the-hell*, following the intuitions developed in the previous section regarding speaker ignorance. This property was initially ascribed to *wh-the-hell* by Oguro (2017), though in his formulation, it simply meant that the speaker did not know the answer to the question, which will be too weak. Here, I argue for a particular flavor of speaker ignorance, not just of the particular answer, but in a sense of *what the answer could even be*. Before that, though, I will (very briefly) introduce the framework for the semantics of questions adopted here.

3.1 Background

I will be largely adopting Kotek's (2018) implementation of Beck's (2006) semantics for questions (see also Beck & Kim 2006; Cable 2010). The idea that the meaning of a question consists of the set of its possible answers (Hamblin 1973) is instantiated in terms of alternatives in the focus-semantic domain. Specifically, *wh*-words contribute no ordinary semantic meaning (are undefined in the ordinary domain), and their focus-semantic meaning is a set of alternatives composed of the possible answers to the question. Thus, all *wh*-words denote Hamblin sets, and *wh*-words differ in the types of alternatives the set they denote contains.

- (16) a. Ordinary value: $\llbracket \text{who} \rrbracket^o$ is undefined (Kotek 2018: 31)
 Alternative value: $\llbracket \text{who} \rrbracket^f = \{x_e : x \in \text{human}\}$
 b. Ordinary value: $\llbracket \text{what} \rrbracket^o$ is undefined
 Alternative value: $\llbracket \text{what} \rrbracket^f = \{x_e : x \notin \text{human}\}$

Other *wh*-words contribute similar restrictions on alternatives, i.e. *when* considers times, *where* locations, etc. The computation of alternatives in the focus domain occurs in parallel to the computation of ordinary semantic values, as proposed by Rooth (1992). To ensure the question has an ordinary semantic value at LF, Kotek introduces the operator ALTS_{HIFT}, 'the source of interrogative semantics' (Kotek 2018: 32), shifting the alternatives of its complement from the focus domain to the ordinary domain. The output of ALTS_{HIFT} is a set of ordinary semantic propositions,

and acts as the input to Dayal (2016)'s answerhood operator ANS, which presupposes that a question has a unique maximally informative true answer, and returns it.

3.2 The proposal: *the hell* as a restricted domain of alternatives

I propose that *the hell* be seen as an additional filter on the alternative sets generated by *wh*-words, roughly one which eliminates all possible answers with which the speaker is familiar. This is implemented through a rather simple semantics, in (17).

$$(17) \quad \llbracket \text{the hell} \rrbracket^f = \bar{\mathbb{C}}$$

A few things to note. First, this gives *the hell* only a focus-semantic value, like any *wh*-word. Second, the focus-semantic value is the complement to a set \mathbb{C} . I define \mathbb{C} as the context set over alternatives, that is, the set of contextually available alternatives. More on that just below, but for now, this means that the denotation of *the hell* is the set of contextually *unavailable* alternatives. This makes it the same type as a *wh*-word, which I will argue is a desirable outcome. So, the composition of a *wh-the-hell* phrase is simply the intersection of the two alternative sets.

$$(18) \quad \llbracket \text{what the hell} \rrbracket^f = \llbracket \text{what} \rrbracket^f \cap \llbracket \text{the hell} \rrbracket^f = \{x_e : x \notin \text{human}\} \cap \bar{\mathbb{C}}$$

So what is \mathbb{C} ? It may be thought of as the set of possible answers to a question which are naturally given as possible in the context, where the context includes common ground both in terms of the discourse and world knowledge. So an answer need not be discourse-familiar to be contextual. \mathbb{C} is thus distinct from, and might be considered the focus-semantic equivalent of, our familiar context set C .

Why is this distinction desirable? Formally, it is necessary to work with the context in the focus-semantic domain, in which case we may want it to be just a type-lifted version of C . But I think that there are more intuitive reasons to treat them as distinct objects: namely, just because you are in the context, doesn't mean I consider you a possible answer. The intuition captured by \mathbb{C} is that the set of individuals who are natural, likely referents in a discourse is not necessarily identical to the set of individuals who are natural, likely answers to a given question. Consider:

(19) Who the hell stole my car?

If you have a dog, odds are that you are quite likely to refer to them, especially at home. But if you utter (19), you are likely not considering *your dog* as a possible answers. The point being, an individual may be in the discourse context as a referent but not in the discourse context as a car thief. When we utter a question like (19), we are not communicating that *all* individuals in the context have been determined to be untrue answers; rather, that all *likely car thieves* seem to be untrue answers.

This notion of \mathbb{C} is admittedly underspecified. How one determines how to divide the space of possible answers between \mathbb{C} and its complement may never be fully answerable in a formal sense. But semantic theory has readily acknowledged the necessity to encode certain features of the context: the discourse participants, the set C of individuals in the context, intensional variables to represent the world, something like the QUD stack (Roberts 2012), and perhaps more. Membership in \mathbb{C} is, I think, no more or less specified than membership in C is.

Before we move on and apply the analysis to capturing *wh-the-hell*'s distribution, note also that this semantics does *not* encode an existential presupposition into *wh-the-hell* questions. (Szabolcsi & Zwarts 1993: 261) argue that there is one, and that *wh-the-hell* questions can only be asked if there is 'unquestionable evidence' that there is a true answer. This seems much too strong, and to me it seems that rhetorical uses are common with *wh-the-hell* questions, especially with modals. The answer *nobody* is the expected one for questions like:

(20) Who the hell would buy a book like that?

Rather, our semantics encodes a local anti-existential presupposition - there may or may not be a true answer, but if there is, it's not one that I am familiar with. From this, we can pragmatically derive the feeling that *wh-the-hell* questions are asked when we strongly expect an answer. Since the question implies no true answer is in the context, it would be strange to ask it if you did not have reason to believe that one existed somewhere. Otherwise, you intend to utter an unanswerable question.

3.3 Wh-less *wh*-questions

You might ask, why give *the hell* the same type as a *wh*-word, rather than treat it as a function taking a *wh*-word as an argument? Motivating this decision is another intriguing property of *the hell*, its ability to license deletion of the *wh*-word.

- (21) a. The hell was that? [COCA, *FBI*, season 1 episode 18]
 b. The hell you doin' in here? [COCA, *The Passage*, season 1 episode 8]

It can also be used on its own, with a meaning equivalent to *what the hell*.

(22) The hell?

In fact, only readings where the elided *wh*-word would be *what* are allowed.

- (23) a. *Someone left. ~~Who~~ the hell was it?
 b. *He left it somewhere. ~~Where~~ the hell was it?

Treating these cases as the only instance of *wh*-ellipsis in the grammar is a tall order, which would be necessary if *the hell* was a function over alternative sets. But (17) lets *the hell* fill the role of a *wh*-word on its own. This also explains the *what*-preference, since nothing in \llbracket the hell \rrbracket triggers an interpretation targeting people or locations. We thus default to the least marked set. Note, however, that it does not bear the [+wh] feature, syntactically, as it cannot be embedded:

(24) *I don't know the hell is going on.

Additionally, this is only possible for the definite non-D-linked constructions, suggesting that perhaps the prepositional ones are, indeed, functions over *wh*-words.

(25) *{On earth/In the hell} is going on?

3.4 Speaker ignorance

(17) implicitly encodes speaker orientation. Contextual availability is dependent on the common ground of the discourse participants, and necessarily relative; available *to me, as the asker*. Thus, ignorance follows from the anti-contextual filtering of *the hell*: any answer the speaker knows is a member of \mathbb{C} . This derives the contrasts in Section 2.1. This section further develops our notion of speaker ignorance.

3.4.1 Absence of perspective shift

The sensitivity that *the hell* shows to the speaker's knowledge state is not logophoric. Some perspectival elements shift whose perspective they represent when embedded in clauses with different attitude holders or perspective takers as subjects. *The hell* is not one of these - it represents the doxastic alternatives of the utterance speaker, even when embedded underneath an alternate attitude holder.

- (26) a. Sam doesn't know who the hell that was. (#But I do.)
 b. Sam believes that they'll catch whoever the hell it was. (#It was me.)

The only instances of shift are embedding under a verb of speech:

- (27) a. Sam asked me who the hell that was, so I told her.
 b. Sam said she didn't know what the hell was going on, but I do.

We might analyze these as a kind of quotation or indirect discourse, which changes the 'speaker' rather than merely the perspective taker. This is problematic for analyses like that of Huang & Ochi (2004) and Akiyama (2014), who posit that *the hell* is sensitive to logophoric attitude or a syntactic point-of-view projection,

respectively. Huang & Ochi (2004) show that Mandarin *daodi* and Japanese *ittai* do undergo this kind of logophoric shift, and argue for a uniform analysis with English *wh-the-hell*. But such uniformity is ill-advised given (26). Similarly, Akiyama (2014) argues that *the hell* enters the derivation unvalued for syntactic POV features, which must be valued by a POV head accessible in the same phase (Chomsky 2008). But this analysis makes incorrect predictions for data like (28):

- (28) Blake knows that Alex won't tell me what the hell is going on.
 [_{POV=Spkr} B knows [_{POV=Blake} that A won't tell me [_{POV=Alex} what...]]]

In (28), *the hell* should be able to be valued by the POV features of its clausal subject Alex, or perhaps the subject of the higher clause, Blake, if *wh-the-hell* is in the phase edge and thus accessible outside of the local clause. But neither of these interpretations are available: the ignorance conveyed by *the hell* is the actual speaker's, despite the POV projection of the speaker being too distant for Agree.

3.4.2 Contextual questions

Anti-contextuality predicts that *the hell* be incompatible with constructions that require a contextually available answer. This seems to hold for a number of examples.

- (29) Since when (*the hell) do you have tenure?

Kiss (2017) says that *Since when p?* can be paraphrased as *Since when is it (or should it have been) part of our common ground that p?* Thus, a *since when* question implies that the proposition being questioned is intended by the addressee to be in the common ground. Given that a specific discourse-salient proposition is necessarily in \mathbb{C} , it is incompatible with *the hell*.

- (30) Guess what (*the hell) I'm doing tonight.

Guess-constructions like (30) have a stronger presupposition: the speaker knows the true answer to the *wh*-question. This is, trivially, incompatible with *the hell*.

3.4.3 Kind questions

Though *the hell* is degraded with NP complements, *kind of* alleviates it.

- (31) What the hell *(kind of) book are you reading?

Akiyama (2014) offers a syntactic explanation, in which he analogizes *kind of* phrases syntactically to measure phrases. But, this does not explain the infelicity without *kind*, and it inaccurately predicts *the hell* to be allowed in measure phrases.

(32) *What the hell {some/half/two} of the books did you read?

The only grammatical equivalents are those which are synonymous.

(33) What the hell {type/sort} of book are you reading?

Thus, the analysis for the contrast in (31) should not be syntactic. I propose a pragmatic account to derive it from *the hell*'s semantics, in the following way. Given that the use of *the hell* in a question indicates that the speaker does not have any contextually available alternatives which they believe could be an answer to the question, replying with a specific individual is uncooperative, as that individual will necessarily be unfamiliar to the speaker, and so will contribute no useful information. Instead, answers which reference properties are more interpretable.

3.4.4 Rogative preference

Recall from 2.1 that *the hell* is compatible with rogatives like *wonder*, but incompatible with *surprise*. This is true even if speaker ignorance is guaranteed.

(34) I don't know who won, but Riley seems surprised by who (*the hell) it is.

So, we cannot assimilate this to our general ignorance requirement. Note instead that *surprise* has a complementary distribution with *wonder* (Karttunen 1977). Only the former can embed a declarative, and only the latter can embed *if*-clauses. *Surprise* + *wh* can only have the reading of being surprised by the answer, and *wonder* can only mean wondering what the answer is. This suggests that *surprise* selects for an answer, while *wonder* selects for the interrogative itself. When the embedding verb is ambiguous, *wh-the-hell* only allows the latter reading, as in (1).

We can explain this via the interaction of ANS and *the hell*. Though ANS can select the maximally informative true answer from \bar{C} , it will still be contextually unavailable - ANS 'knows' it exists, but cannot make it discourse-available. *Surprise*, which requires an available answer, will be unable to select such an output. *Wonder* doesn't need the answer. We could say that rogatives embed a question without ANS, selecting directly for a Hamblin set, but I will not develop this idea further here.

3.4.5 Ignorance implications

While the general claim from DG that negative implicatures license *wh-the-hell* is too strong, this notion is useful to explain another pattern. Specifically, predicates vary on the kind of ignorance implication that they generate out of the blue, and this corresponds to a spectrum of acceptability with embedded *wh-the-hell*.

(35) a. I {found out/discovered} who the hell stole my car.

- b. ?I {confirmed/realized} who the hell stole my car.
- c. *I {heard/was told} who the hell stole my car.

(35a) strongly implies the subject did not know the embedded proposition before, and underwent some search to find the answer. (35b) does so to a lesser extent, and (35c) has no implication of past ignorance. These judgments, however, are sensitive to contextual manipulation. Observe that they can all embed *wh-the-hell* if they are modified by *finally*, which provides the necessary ignorance implication by itself.

3.5 Defending domain restriction

The main thrust of this analysis is that *wh-the-hell* should be seen as restricting the *wh*-word's domain of quantification, rather than expanding it. This contrasts with DG's approach, which says that *the hell* 'extends the domain of quantification' of the *wh*-word 'to include familiar and novel values', and as a result, 'the domain of quantification for *wh-the-hell* is the entire domain D, and not just a presupposed subset of it, as with regular *wh*-words.' (DG: 43). Then, a scalar inference is triggered such that *wh-the-hell* implies the answer is in the extended domain. This argument is independent from the polarity concerns, so it needs its own treatment.

3.5.1 The *wh*-typology

My primary argument is this: treating *wh-the-hell* as domain expansion allows a contrast between it and *what*, but not between *what* and *which*. It undesirably collapses the typology of *wh*-questions and strands those *wh*-words which are neither aggressively D-linked nor non-D-linked. The scalar inference DG describe needs to run between *what* and *which* to derive their behavior: *what* quantifies over the whole domain, and *which* over the local subdomain, so the use of *what* implies a less familiar answer. Adding *the hell*, instead, restricts *what* to the non-local domain, which also predicts **which the hell*, something not derived by the expansion account.

This is required for when they both take an NP complement. Imagine that a professor has assigned students to each read a paper, but let them choose which one. If the professor provides a restricted list, (36b) is preferred and (36a) is marked. With no relevant list, the inverse holds. If both options are available, from a list or your own choice, (36b) implies you used the list and (36a) the opposite.

- (36) a. What paper did you choose?
- b. Which paper did you choose?
- c. What the hell did you choose?

In all contexts, (36c) not just implies, but requires the speaker believe the addressee chose an unknown paper. This pattern cannot be derived by the expansion account of *wh-the-hell*. Additionally, it motivates the inclusion of a contextual restriction in $\llbracket \text{which} \rrbracket$. Deriving the contextual nature of *which* arises from its NP complement fails to predict the divergent behavior of *what NP* and *which NP*.

3.5.2 So-called D-linked uses

An obvious problem for the domain restriction account, on the other hand, is the existence of evidently D-linked uses of *wh-the-hell*:

- (37) What the hell is the right answer? (to a multiple choice question) (Oguro 2017: 117)
- (38) Who on earth will they pick? (in a reality show) (Rawlins 2008: 196)

In these examples, there is a defined list of possible answers. However, I would argue that they are only licit *if the speaker has reason to reject every answer*. (37) is only acceptable if you have already, for example, plugged each answer into the formula and gotten the wrong output every time; it would be infelicitous said out of the blue upon encountering the question. Similarly, (38) is only acceptable if no contestant is a viable winner. While Rawlins compares it to *Whoever will they pick?*, he also acknowledges that ‘In the case of the expletive question, the speaker might also be suggesting that it seems implausible that they could choose anyone at all’. This is the key difference: *-ever* suggests ignorance between the possible answers, which is possible if they have all performed well as in the context he provides, while *the hell* suggests that none seem viable. In this way, using *wh-the-hell* in a supposedly D-linked context is a pragmatic move to declare your lack of belief in the provided answer set, while *-ever* only suggests you cannot decide between them.

Additionally, Rawlins’ intensional domain widening, while more persuasive than DG’s extensional widening, still fails to make certain relevant predictions. In the next sections, I will argue that it is the fact that *wh-the-hell* actually *excludes* the context that explains its behavior with *which*, non-fronted questions, and sluices.

3.6 *Which the hell

Based on the arguments above, I revise Kotek’s semantics for *which* to include an explicit contextual restriction, rather than assuming it is provided by the NP.

$$(39) \quad \llbracket \text{which} \rrbracket^f = \lambda P_{\langle e,t \rangle}. \lambda x_e. P(x) \wedge x \in \mathbb{C} = \lambda P_{\langle e,t \rangle}. P \cap \mathbb{C}$$

We could instead stipulate that *which* only combines with NPs carrying a contextual restriction, but that feels like quite the same move. Now, the set denoted by

which NP is a subset of the set denoted by *what NP*. This allows the scalar inferences as described in Section 3.5 to proceed as intended. The incompatibility of *which* and *the hell* should also now be obvious. Since *which* filters out alternatives not in \mathbb{C} , and *the hell* filters out alternatives in \mathbb{C} , the result will be the empty set.

$$(40) \quad \llbracket \text{which the hell} \rrbracket^f = \llbracket \text{which} \rrbracket^f \cap \llbracket \text{the hell} \rrbracket^f = \mathbb{C} \cap \bar{\mathbb{C}} = \emptyset$$

The derivation will then crash at ANS. Another consequence of this discussion is that *which*, and thus \mathbb{C} , imposes a looser restriction than the discourse familiarity or salience requirement of a definite. This is desirable, given that (41a) is licit in contexts that (41b) is not: only the latter requires a salient subset of countries.

- (41) a. Which countries did you visit?
 b. Which of the countries did you visit?

3.7 Single in-situ questions

Recall that *wh-the-hell* cannot be in-situ, as in (3d), or echoed, as in (2). Both of these can be derived from the anticontextuality requirement, at least for single *wh*-questions; we'll return to multiple questions in Section 4.

3.7.1 Non-echo questions

Pires & Taylor (2007) and Bobaljik & Wurmbrand (2015) both observe that single *wh*-questions may occur non-fronted even when not echoed in English. Such constructions are felicitous, it seems, only when there is the assumption that the listener knows the answer already.

- (42) a. I was driving along Andrews Avenue. (Pires & Taylor 2007)
 b. And you were driving in which direction?
- (43) a. So, your boy's name is what? [McNulty, *The Wire*, S1E1]
 b. You've been in this post how long, exactly? [*Harry Potter and the Order of the Phoenix*, Ch. 15]

Bobaljik & Wurmbrand (2015) argue that such questions lack interrogative syntax, as they cannot embed under the same predicates as fronted *wh*-questions. In the absence of a clausal $C_{[+Wh]}$, then, the in-situ *wh*-phrase carries a phrasal *Q* feature. Interrogative semantics is triggered pragmatically: the *wh*-word, like an indefinite, triggers a set of alternatives. Unlike an indefinite, however, the *wh*-word contributes no ordinary value, and so it acts as a sort of placeholder in the sentence, prompting the listener to choose between the proffered alternatives. This pragmatic reasoning only functions if the speaker can assume the listener knows the true answer, and thus implies a contextuality requirement incompatible with *the hell*.

3.7.2 Echo questions

Beck & Reis (2018) give a compelling analysis of echo questions on which there is also no clausal interrogative syntax, the echo *wh*-phrase has a phrasal Q, and it is the narrow focus on the *wh*-word which triggers interrogative semantics: ‘Focus marking... triggers the introduction of alternatives... In the case of narrow focus on *wh*, there is – unusually – one particular alternative, namely the deictic or anaphoric alternative to the *wh*-element... Focus, as always, requires a (adjacent) discourse antecedent. Thus the deictic/anaphoric alternative has to be available in the context.’ (Beck & Reis 2018: 16) Thus, the meaning of an echoed *wh*-word is as in (44), and this results in a clash with *the hell* as in (45), giving an unanswerable question.

(44) $\llbracket \text{WHAT} \rrbracket^f = \{z\}$ (where z is the unique contextually relevant element of D_e)

(45) $\llbracket \text{WHAT the hell} \rrbracket^f = \llbracket \text{WHAT} \rrbracket^f \cap \llbracket \text{the hell} \rrbracket^f = \{z\} \cap \overline{\mathbb{C}} = \emptyset$

3.8 Sluicing

Since Merchant (2001), it has been noted that *wh-the-hell* cannot be sluiced.

(46) Someone bought this book, but I don’t know who (*the hell).

DG attribute this to an inability to establish an anaphoric link between *wh-the-hell* and the discourse-familiar *someone*, following from the assumption that *wh-the-hell* does not carry an index. But this analysis cannot be correct, given that *wh-the-hell* can be linked anaphorically to an indefinite or definite NP:

(47) a. Someone_i bought that book. John won’t tell me who_i the hell it was.

b. I’ve trailed the thief_i for weeks and I still don’t know who_i the hell he is.

Thus, we need an independent account of the sluicing fact. Syntactic identity accounts like Rudin (2019) will not do, since the elided material in (46) is the same with or without *the hell*. The semantic identity account in Merchant (2001), which requires symmetric entailment between the elided material and its antecedent, will also fail, since the truth conditions of the elided material are also the same as the antecedent, in virtue of *the hell* contributing no ordinary meaning.

We can find what we need in the inquisitive semantic account of AnderBois (2014), who notes that the only acceptable antecedents for sluicing are those which are inquisitive (introduce alternatives). Observing that sluicing is not licensed where two clauses are truth-conditionally identical but introduce different alternatives, AnderBois revises Merchant’s symmetric entailment condition in two ways. First, entailment is defined over alternatives, mirroring Groenendijk & Stokhof (1984):

(48) $\phi \models \psi$ iff $\forall \alpha \in \llbracket \phi \rrbracket$, α is s.t. $\exists \beta \in \llbracket \psi \rrbracket$ s.t. $\alpha \subseteq \beta$ (AnderBois 2014: 899)

The second revision is that the symmetric entailment condition holds not of the elided IP and its antecedent IP, but of the CPs containing them.

(49) SYMMETRIC ENTAILMENT CONDITION ON SLUICING: Given a structure $[CP_E[C + QIP_E]]$, IP_E can be elided only if there is some salient antecedent CP_A such that $CP_E \models CP_A$, and $CP_A \models CP_E$. (AnderBois 2014: 900)

Combined with (17), this straightforwardly derives the sluicing ban. Consider (46). Symmetric entailment is trivial with *who* condition, as the alternatives of *who* and *someone* are identical (Kratzer & Shimoyama 2002). Thus, since $\llbracket \text{who} \rrbracket^f = \llbracket \text{someone} \rrbracket^f$, (48) will be trivially satisfied, and since the rest of the clause is identical, (49) will be as well. For each alternative α in $\llbracket \text{someone bought this book} \rrbracket$, an alternative $\beta = \alpha$ exists in $\llbracket \text{who bought this book} \rrbracket$.

This is not the case with *who the hell*. As we saw before, while $\llbracket \text{who} \rrbracket^f$ is contextually unconstrained, $\llbracket \text{who the hell} \rrbracket^f = \llbracket \text{who} \rrbracket^f \cap \overline{\mathbb{C}}$, such that the set of alternatives introduced by *who the hell* is strictly a subset of the set introduced by *who*, and thus of the set introduced by *someone*. Since $\llbracket \text{someone} \rrbracket^f$ includes alternatives in both \mathbb{C} and $\overline{\mathbb{C}}$, there will necessarily be alternatives it introduces which do not have a corresponding alternative in $\llbracket \text{who the hell} \rrbracket^f$, which contains solely alternatives in $\overline{\mathbb{C}}$. Formally, $\forall \alpha \in \llbracket \text{someone} \rrbracket^f$ [$(\alpha \in \mathbb{C}) \rightarrow \neg \exists \beta \in \llbracket \text{who the hell} \rrbracket^f$ s.t. $\alpha \subseteq \beta$]. Thus, (49) will never be satisfied in this case where CP_E is a subset of CP_A , since $CP_E \models CP_A$ will always be true but $CP_A \models CP_E$ will never be.

No account of sluicing not referencing alternatives will be able to capture (46).

4 D-linking and multiple in-situ

Finally, this section considers the ban on *the hell* in the lower position of a multiple *wh*-question, as in (3c), and proposes assimilating it to a more general restriction. To set up our discussion of multiple questions, let's look at one more illustrative pattern. Certain adverbs of quantity and frequency are incompatible with *wh-the-hell*.

- (50) a. *I {largely/mostly/usually/rarely/often} don't know who the hell cheats.
 b. *For the most part/With few exceptions, I don't know who the hell cheats.

But no such problem with *never*, or with *always* outside the scope of negation:

- (51) a. I never know who the hell cheated.
 b. I always find out who the hell cheated.
 c. *I don't always know who the hell cheated.

Per Lahiri (1991), such adverbs, modifying question-embedding predicates, quantify over parts of the complete answer to the question, e.g. *mostly* means that you know most of the parts of the answer. The relevant insight is that they require quantification over individual sub-answers composing a larger answer, where such sub-answers will be individual alternatives. Since the alternatives introduced by *wh-the-hell* are unavailable to the speaker, the ability to access individual sub-parts of the complete answer will be restricted. So, the acceptable forms of quantification address the entire set at once - (51a) and (51b), both universal quantification - but as soon as you need to 'look into' the set and accept or reject specific parts, those parts would need to be available to you. So, existential quantification like (51c) or *most*-type quantification like (50) cannot occur over *wh-the-hell*.

The intuition that we are unable to evaluate individual possible answers with *wh-the-hell* is, I think, critical to what D-linking actually is. I want to adopt an idea mentioned in passing by Comorovski, that 'the use of D-linked *wh*-phrases requires that the participants in the conversation be able to partition identically the set that the *wh*-phrase ranges over' (Comorovski 1996: 13). This is partitioning not in the G&S sense, but rather the mutual division of the set of possible answers by speaker and hearer into pieces of equal size. This identical partitioning requirement is trivial with contextually available sets, like *which NP*, since the maximum partition into singletons is always possible with known members. Conversely, *wh-the-hell* can never be D-linked in this sense, as the subparts of \bar{C} are unknown.

D-linking on this view is a gradient property of the context, not a binary feature of *wh*-words. Different *wh*-words will be differentially *D-linkable*, by how naturally discrete the answers they range over are, on a spectrum between *which* and *wh-the-hell*. The individuals ranged over by *who* are easily partitioned, and treated as atomic; locations in *where* are generally discrete; and the sets of reasons and manners denoted by *why* and *how* will be difficult to D-link absent rich abnormal context. 'Aggressive' non-D-linking, then, is just the extreme case, a semantics which *precludes* such individuation; and the inverse for *which*, which guarantees it.

What I want to draw attention to is that this spectrum correlates with in-situ acceptability. While *which* can always be in-situ, and *wh-the-hell* never is, other *wh*-words vary based on D-linkability. To illustrate, imagine two versions of a cooking competition. In one, cooks prepared dishes which are now lined up anonymously in front of judges, and after evaluation, the judges ask the cooks to come forward and stand next to their dish. In another, it is the contestants who are lined up and the dishes hidden, and judges ask each cook to retrieve and present their dish. Only in the former context, where the dishes are known, is (52) an acceptable question.

(52) Who made what?

Similarly, in-situ *where* is gradiently improved with more identifiable answers.

- (53) *Asking colleagues about their summer travel destinations: Who went where?*
a. ✓ Students went to one of five designated countries for study abroad.
b. ? Linguists conducted fieldwork in potentially any country in the world.
c. # NASA travelers documenting previously unknown parts of the universe.

Even *why* and *how*, claimed to be impossible in-situ because they are Merged high (Rizzi 1990; Stepanov & Tsai 2008), can be tolerated if the context provides a clear individuation of the possible answers.

- (54) *Contestants restricted to one of three specified methods for cooking, selected ahead of the challenge: ?Who cooked their dish how?*
(55) *Contestants divided into teams based on their choice of one of three provided reasons: ?Who is competing why?*

All I can do here is suggest that this pattern may indeed be generalizable, but if true, it allows us to assimilate the ban on in-situ *wh-the-hell* to a more general one. But why? If non-D-linked *wh*-phrases can generally not be in-situ, our explanation cannot reference a specific feature of *wh-the-hell*, and reducing the *wh-the-hell* problem to a general one is all that is in our scope for this paper about *wh-the-hell*. So, more questions than answers. But Kotek (2018) offers us a potential out, arguing based on focus intervention effects that no in-situ *wh*-phrases undergo covert movement to CP, but rather are interpreted via a long-distance focus computation, composing alternatives via pointwise FA. With limited space, here is the briefest of gestures towards a proposal: at-a-distance FA in the focus domain, being *pointwise*, may require such individuated alternatives in Comorovski's sense. This has immediate intrigue and problems; consider it an open invitation to see where it goes.

5 Conclusion

I have offered an account of *wh-the-hell* as a restricted domain of anti-contextual alternatives, taking seriously the idea of non-D-linking as its fundamental property, which is empirically superior to a polarity-based account. It predicts *wh-the-hell*'s embedding behavior, its restrictions on *wh*-positions, and its inability to be sluiced, thus arguing in favor of an inquisitive theory of sluicing. It also introduces a new notion of a question-specific context. It proposes assimilating the ban on *wh-the-hell* in-situ to a general ban on non-D-linked phrases, and opens up the question of how context, focus alternatives, and *wh*-movement rules interact to generate such restrictions. In doing so, I hope to have shown that *wh-the-hell* is an illuminating test case for interrogative syntax and semantics, rather than a fringe phenomenon.

References

- Akiyama, Masahiro. 2014. A formal characterization of *the hell* as a marker of bafflement and phase impenetrability. *Sophia Linguistica* 61. 1.
- AnderBois, Scott. 2014. The semantics of sluicing: beyond truth conditions. *Language* 90(4). 887–926. doi:10.1353/lan.2014.0110.
- Beck, Sigrid. 2006. Intervention effects follow from focus interpretation. *Natural Language Semantics* 4. 1–56. doi:10.1007/s11050-005-4532-y.
- Beck, Sigrid & Shin-Sook Kim. 2006. Intervention effects in alternative questions. *Journal of Comparative Germanic Linguistics* 9. 165–208. doi:10.1007/s10828-006-9005-2.
- Beck, Sigrid & Marga Reis. 2018. On the Form and Interpretation of Echo *Wh*-Questions. *Journal of Semantics* 35(3). 369–408. doi:10.1093/jos/ffy002.
- Bobaljik, Jonathan & Susi Wurmbrand. 2015. Questions with Declarative Syntax Tell us What about Selection? In Ángel J. Gallego & Dennis Ott (eds.), *50 Years Later: Reflections on Chomsky's Aspects*, vol. 77 MIT Working Papers in Linguistics, 13–31. Cambridge, MA: MIT.
- Cable, Seth. 2010. *The grammar of Q: Q-particles, wh-movement, and pied-piping*. Oxford: Oxford University Press.
- Chomsky, Noam. 2008. On phases. In Robert Freidlin, Carlos P. Otero & Maria Luisa Zubizarreta (eds.), *Foundational issues in linguistic theory: Essays in honor of Jean-Roger Vergnaud*, 133–167. Cambridge, MA: MIT Press.
- Comorovski, Ileana. 1996. *Interrogative Phrases and the Syntax-Semantics Interface*, vol. 59 Studies in Language & Philosophy. Dordrecht: Springer Netherlands.
- Dayal, Veneeta. 2016. *Questions* Oxford Surveys in Semantics and Pragmatics. Oxford: Oxford University Press.
- den Dikken, Marcel & Anastasia Giannakidou. 2002. From Hell to Polarity: "Aggressively Non-D-Linked" *Wh*-Phrases as Polarity Items. *Linguistic Inquiry* 33(1). 31–61. doi:10.1162/002438902317382170.
- Giannakidou, Anastasia. 1999. Affective dependencies. *Linguistics and Philosophy* 22. 367–421. doi:10.1023/A:1005492130684.
- Groenendijk, Jeroen & Martin Stokhof. 1984. *Studies on the semantics of questions and the pragmatics of answers*. Amsterdam, Netherlands: University of Amsterdam PhD dissertation.
- Hamblin, Charles. 1973. Questions in Montague English. *Foundations of Language* 10. 41–53. doi:10.1016/B978-0-12-545850-4.50014-5.
- Huang, C.-T. James & Masao Ochi. 2004. Syntax of the hell: Two types of dependencies. In Moulton Keir & Matthew Wolf (eds.), *34th Conference of the North Eastern Linguistic Society*, 279–293.
- Karttunen, Lauri. 1977. Syntax and semantics of questions. *Linguistics and Philos-*

- ophy 1. 3–44. doi:10.1007/BF00351935.
- Kiss, Angelika. 2017. Meta-conversational *since when*-questions and the common ground. In Volha Petukhova & Ye Tian (eds.), *21st Workshop on the Semantics and Pragmatics of Dialogue*, 38–47. Saarbrücken: Saarland University.
- Kotek, Hadas. 2018. *Composing questions*, vol. 80 Linguistic Inquiry Monographs. Cambridge, MA: MIT Press.
- Kratzer, Angelika & Junko Shimoyama. 2002. Indeterminate pronouns: the view from Japanese. In Y. Otsu (ed.), *3rd Tokyo Conference on Psycholinguistics*, 1–25. Tokyo: Hituzi Syobo.
- Lahiri, Utpal. 1991. *Embedded interrogatives and predicates that embed them*: MIT PhD dissertation.
- Merchant, Jason. 2001. *The Syntax of Silence* (Oxford Studies in Theoretical Linguistics 1). Oxford University Press.
- Obenauer, Hans-Georg. 1994. *Aspects de la syntaxe a-barre: Effets d'intervention et mouvements des quantificateurs*: Universite de Paris VIII PhD dissertation.
- Oguro, Takeshi. 2017. The true nature of wh-the-hell phrases. In Stefan Huber & Sonia Ramírez Wolmuth (eds.), *2nd Annual Tampa Workshop on Syntax, Semantics, and Phonology*, vol. 4 (Tampa Papers in Linguistics 2), .
- Pesetsky, David. 1987. Wh-in-situ: movement and unselective binding. In Eric J. Reuland & Alice G.B. Ter Meulen (eds.), *The representation of (in)definiteness*, vol. 14 Current Studies in Linguistics, 98–129. Cambridge, MA: MIT Press.
- Pires, Acrisio & Heather Lee Taylor. 2007. The Syntax of Wh-In-Situ and Common Ground. In *Annual Meeting of the Chicago Linguistic Society*, vol. 43 2, 201–215.
- Rawlins, Kyle. 2008. *(Un)conditionals: An investigation in the syntax and semantics of conditional structures*: UC Santa Cruz PhD dissertation.
- Rizzi, Luigi. 1990. *Relativized Minimality*. Cambridge, MA: MIT Press.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics* 5(6). 1–69. doi:http://dx.doi.org/10.3765/sp.5.6.
- Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1(1). 75–116. doi:10.1007/BF02342617.
- Rudin, Deniz. 2019. Head-Based Syntactic Identity in Sluicing. *Linguistic Inquiry* 50(2). 253–283. doi:10.1162/ling_a_00308.
- Stepanov, Arthur & Wei-Tien Dylan Tsai. 2008. Cartography and licensing of *wh*-adjuncts: a cross-linguistic perspective. *Natural Language and Linguistic Theory* 26. 589–638. doi:10.1007/s11049-008-9047-z.
- Szabolcsi, Anna & Frans Zwarts. 1993. Weak islands and an algebraic semantics of scope taking. *Natural Language Semantics* 1(3). 235–284. doi:10.1007/BF00263545.

Joshua Martin
Harvard University Department of Linguistics
Boylston Hall, Cambridge MA 02138
joshuamartin@g.harvard.edu