In: Eszter Ronai, Laura Stigliano and Yenan Sun (eds.). 2019. *Proceedings of the Fifty-Fourth Annual Meeting of the Chicago Linguistic Society*. pp.451-460.

The Acquaintance inference with *seem*-reports*

Rachel Etta Rudolph
University of California, Berkeley

1 Introduction

Some assertions give rise to the *acquaintance inference*: the inference that the speaker is acquainted with some individual (Ninan 2014). Discussion of the acquaintance inference has previously focused on assertions about aesthetic matters and personal tastes, as in (1); but it also arises with reports about how things *seem*, as in (2).

(1) a. **A:** The cake is tasty. ⇝ A has tasted the cake.
b. # **A:** The cake is tasty, but I haven't tasted it.

(2) a. **A:** Tom seems like he's cooking. ⇝ A has perceived Tom.
b. # **A:** Tom seems like he's cooking, but I haven't perceived him.

The acquaintance inference with 'seem'-reports, though acknowledged (Pearson 2013, Ninan 2014), has not yet been adequately explored.[1] Reports with 'seem' give rise to puzzling acquaintance behavior, with no analogue in the previously-discussed domains (§2). In particular, these reports call for a distinction between the specific acquaintance inference (that the speaker is acquainted with a specific individual) and the general acquaintance inference (that the speaker is acquainted with something or other of relevance). We frame a novel empirical generalization — the specific with stage-level generalization — that systematizes the observed behavior, in terms of the semantics of the embedded 'like'-clause (§3). To preview: we hold that which type of acquaintance inference a 'seem'-report gives rise to depends on whether the embedded clause contains a stage-level or an individual-level predicate. We present supporting experimental work (§4), and explain why the generalization makes sense given the evidential role of 'seem'-reports (§5). Finally, we discuss the relevance of this result for extant proposals about the semantics of 'seem'-reports (§6). More modestly, it fills a gap in previous theories by identifying which reports get which of two possible interpretations; more radically, it suggests a revision of the kind of explanation that should be given for the acquaintance behavior in question.

[0]Many thanks to the organizers of CLS 54, Eszter Rónai, Laura Stigliano, and Yenan Sun. This paper benefited from feedback from participants at CLS, as well as at the Berkeley Research Group in Formal Semantics, CUSP 10 at UC Irvine, the PhLing Workshop at Northwestern, and the Linguistics and Philosophy Workshop at the University of Chicago. For helpful conversation, feedback, and guidance, thanks to Amy Rose Deal, David Gottlieb, Chris Kennedy, Tania Lombrozo, John MacFarlane, Patrick Muñoz, Joe Roussos, Tamara Vardomskaya, Malte Willer, Ming Xiang, and Seth Yalcin. All errors are my own.

[1]While this paper solely addresses reports with 'seem', there is also a lot to say about reports of similar forms with the specific sensory appearance verbs ('look', 'smell', etc.). See (Rudolph 2019) for discussion of this broader class.

2 Two types of acquaintance inference

The 'seem'-report in (2a), repeated in (3), is a *copy raising* (CR) construction.

(3) **A:** Tom seems like he's cooking. $\rightsquigarrow$ A has perceived Tom.

Copy raising constructions are characterized by having a substantive subject and an embedded 'like'-clause containing a pronoun that corefers with it. A contrast has been observed between CR reports and their *expletive subject* (ES) variants, as in (4) (Asudeh & Toivonen 2012, Rogers 1972).[2]

(4) **A:** It seems like Tom is cooking.
$\rightsquigarrow$ A has perceived something relevant to whether Tom is cooking.

If A doesn't see Tom, but sees his kitchen with preparations for dinner apparently underway, the ES report in (4) is appropriate, but the CR report in (3) is not. To characterize this situation, we introduce terms to distinguish two types of acquaintance inference: first, the *specific acquaintance inference* that the speaker is acquainted with a specific individual; second, the *general acquaintance inference* that the speaker is acquainted with something or other that is evidentially-relevant. Assertions about aesthetic matters and personal tastes, as in (1), as well as the CR report in (3), give rise to the specific acquaintance inference. To be appropriate, they require the speaker to have first-hand acquaintance with the specific individual named by the subject. The ES report in (4), by contrast, only gives rise to the general acquaintance inference. It is appropriate so long as the speaker has perceived something that is relevant to the question of whether Tom is cooking.[3]

The previous two examples may suggest that copy raising reports all have the specific acquaintance inference, while expletive subject reports have the general. But this, admittedly tempting, generalization would be too quick. Some CR reports also only give rise to the general acquaintance inference. If A walks into Tom's kitchen and notices vegetables partially chopped on the cutting board, all exactly even, and a perfectly-cooked roast cooling on the counter, they can make either report in (5), even if Tom isn't present.

(5) a. **A:** It seems like Tom is an experienced cook. ES
b. **A:** Tom seems like he's an experienced cook. CR

Why do some copy raising 'seem'-reports give rise to the specific inference, and some only to the general? This is the central question we address in this paper.

3 The specific with stage-level generalization

There are likely many factors that influence the acquaintance behavior of copy raising 'seem'-reports. Here, we offer evidence for one novel generalization:

[2] 'Seem'-reports can also take other forms. The discussion here is restricted primarily to copy raising and expletive subject constructions, though we also touch briefly on "copy free" reports in §6.

[3] Under the plausible assumption that the individual the embedded claim is about is always relevant to the truth of that claim, the presence of the specific inference entails the presence of the general one. The reverse, of course, is not the case.

Specific with stage-level generalization CR 'seem'-reports with embedded *stage-level predicates* (SLPs) give rise to the specific acquaintance inference; those with embedded *individual-level predicates* (ILPs) give rise only to the general acquaintance inference.

SLPs are predicates which apply primarily to stages of individuals, and are presumed to hold transiently. Examples include 'cooking' and 'drunk'. ILPs, by contrast, are predicates which apply to individuals, and are presumed to hold in a more standing way. Examples include 'an experienced cook' and 'well-organized'.[4] We present two diagnostics for this distinction.

Bare plural diagnostic Bare plural subjects of SLPs have existential interpretations, while bare plural subjects of ILPs have universal or generic interpretations.

(6) a. Students are cooking. ∃: SLP
b. Students are upset. ∃: SLP
c. Students are experienced cooks. ∀: ILP
d. Students are well-organized. ∀: ILP

Absolute construction diagnostic (Stump 1985) Absolute constructions with SLPs in the antecedent are equivalent to conditionals with 'if…' or 'when…', while those with ILPs are equivalent to conditionals with 'since…' or 'given that…'.

(7) a. Cooking, Tom is happy. 'if…': SLP
b. Upset, Tom likes to be alone. 'if…': SLP
c. An experienced cook, Tom prefers to eat in. 'since…': ILP
d. Well-organized, Tom is good at his job. 'since…': ILP

So, to repeat, the specific with stage-level generalization states, first, that copy raising 'seem'-reports with embedded stage-level predicates, as in (3) with 'cooking', give rise to the specific acquaintance inference that the speaker has perceived the individual denoted by the subject; and second, that copy raising 'seem'-reports with embedded individual-level predicates, as in (5b) with 'an experienced cook', only give rise to the general acquaintance inference that the speaker has perceived something or other relevant to the truth of the embedded claim.

4 Experimental work

The specific with stage-level generalization is supported by new experimental work comparing six minimal ES/CR pairs — three with embedded SLPs ('cooking', 'upset', 'playing outside'), and three with embedded ILPs ('an experienced cook', 'well-organized', 'enjoys arts and crafts'). Subjects recruited through Amazon Mechanical Turk ($N = 719$) were presented with short scenarios, and asked to rate the correctness of an utterance in the given scenario on a scale from 1 to 7 (see Fig. 1). Each subject saw only one test sentence (in addition to one filler question, and two practice questions that also served as attention checks).

[4] The distinction goes back to Carlson (1977).

The ES and CR reports were tested in scenarios where the speaker did not have perceptual acquaintance with the individual denoted by the matrix CR subject. The specific with stage-level generalization then makes the following predictions. First, with the pairs with embedded SLPs (see sample stimulus in Fig. 1), there would be a significant difference in acceptability between the two report types, with the ES reports being rated higher than the CR reports. Second, with the pairs with embedded ILPs (see Fig. 2), there would be no such difference: both CR and ES pairs would be ranked equally, even though the CR matrix subject was not perceived.

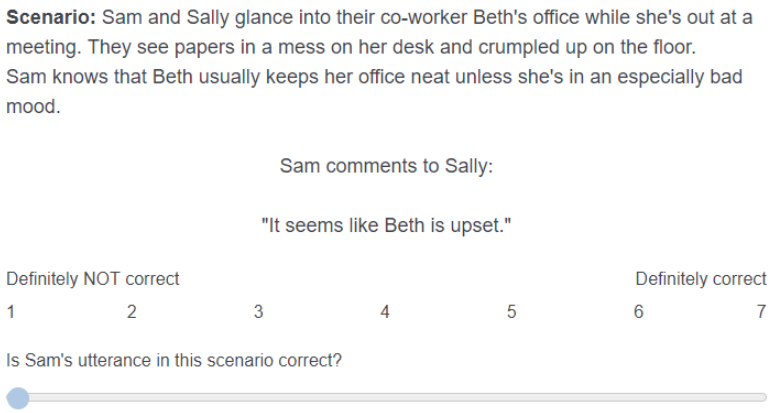

Figure 1: Sample experimental stimulus: SLP (ES version)

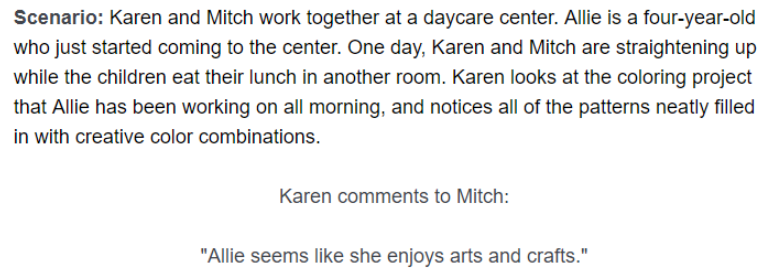

Figure 2: Sample experimental stimulus: ILP (CR version)

Results were in line with the predictions of the specific with stage-level generalization. In the SLP cases there was a significant effect of report type on the ratings of utterance correctness;[5] in the ILP cases, there was no such effect (Fig. 3).[6]

[5]The results of one-way ANOVA tests are as follows. For the pair with 'cooking': $F(1, 111) = 14.81, p < .001$ (as homogeneity of variance could not be assumed, we also report the Welch statistic: $F(1, 72.4) = 12.89, p = .001$); 'upset': $F(1, 103.43) = 5.41, p = .02$ (Welch); 'playing outside': $F(1, 116) = 13.22, p < .001$. The smaller effect with 'upset' suggests that stative SLPs may tolerate a reading with only the general acquaintance inference more easily than eventive SLPs. Further tests will be needed to see if this generalization holds up.

[6]For the pair with 'an experienced cook': $F(1, 131) = .91, p = .34$; 'well-organized':

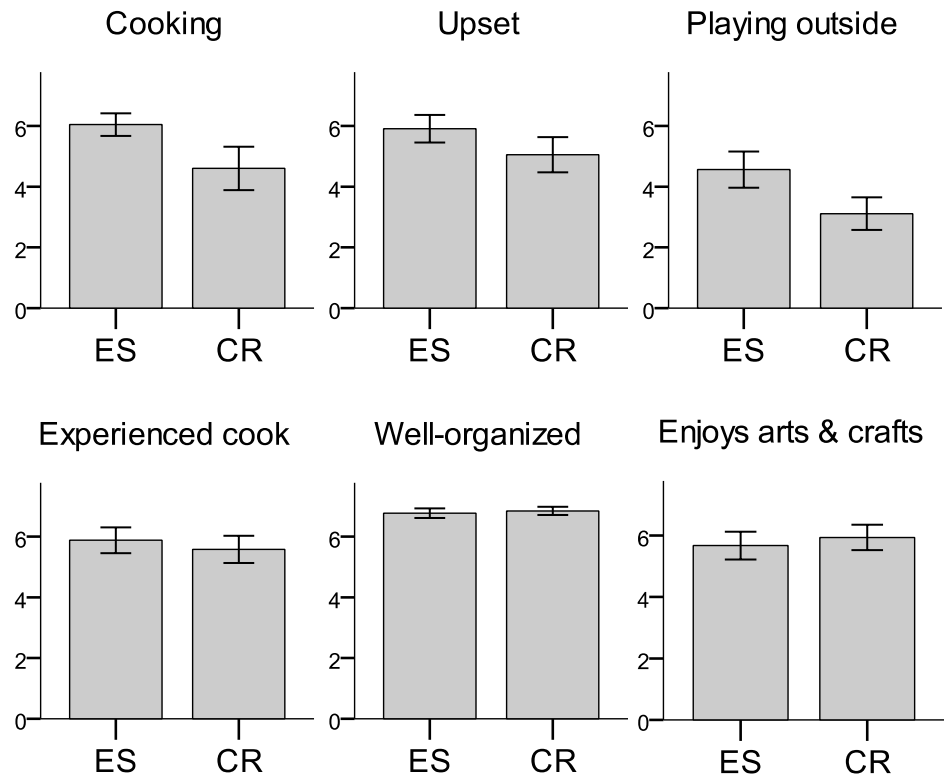

Figure 3: Mean rankings of utterance correctness (top: SLPs, bottom: ILPs)

5 Why do stage- and individual-level predicates behave differently?

The specific with stage-level generalization can be given a functional rationale. 'Seem'-reports have an evidential role: they are used to express that one has perceptual evidence for some state of affairs, specified in the embedded clause. Since SLPs denote properties presumed to hold more transiently, perception of the individual the claim is about tends to be better, more direct evidence than perception of some scene not containing that individual. For example, perceptual evidence for someone cooking or being upset will generally be better if it is perception of that person; perception of some scene that may suggest that the relevant state of affairs holds, but doesn't include the individual, will tend to be worse. Thus, it makes sense for reports embedding SLPs to have a way to mark this distinction in evidential situation. Reserving CR reports for the special case of direct perception of the target individual would do this.

ILPs, by contrast, denote properties that individuals have in a more standing way. Because of this, there is less of a contrast between the evidential power of perception of that individual, compared with perception of some other scene.

$F(1, 124) = .46, p = .5$; 'enjoys arts and crafts': $F(1, 116) = .75, p = .39$.

For example, if the question is whether someone is an experienced cook, or well-organized, it matters less what they look like at the moment. Perception of some scene appropriately related to them can be just as good, or even better evidence that the embedded claim obtains. Thus, reports with embedded ILPs don't call for the same contrast between the two kinds of perceptual evidential situations the speaker may be in, and so don't reserve CR reports for just one of them.[7,8]

This, then, is a sketch of a functional basis for the differing acquaintance behavior of 'seem'-reports embedding SLPs and ILPs. In the next section, we turn to a discussion of how this behavior might be captured within a semantic analysis of copy raising and expletive subject constructions.

6 Theoretical discussion

Recent literature on the semantics of copy raising constructions addresses the question of whether they assign the role of *perceptual source* (p-source) to the matrix subject. Asudeh & Toivonen (2012) characterize the perceptual source as "what is perceived in a perceptual event or state" (p. 322). They take the unacceptability of the copy raising report from (3), repeated here in (8a), in the given "absent cook" scenario, to be grounds for taking the matrix subject, 'Tom', to be assigned the role of p-source — as shown in (8b). And they take the acceptability of the expletive subject variant in (8c), in the same scenario, to show that no particular constituent in that report is assigned the p-source role. Instead, they take the role to be existentially quantified over, as in (8d).[9]

---

[7]Our claim here about copy raising reports embedding SLPs is similar to what Rett & Hyams (2014) say about all CR reports: holding that they encode "direct evidentiality," while ES reports do not — where the direct evidential component is that the speaker has perceived the referent of the matrix subject (see esp. pp. 176–179). We depart from them in holding that we must recognize variability across CR reports. This variability could be captured by denying that direct evidentiality is communicated by CR reports embedding ILPs. Alternatively, it could be captured by holding that there is variability in what counts as "direct evidence": that is, perhaps for claims with ILPs, direct evidence doesn't have to include the target individual. (These options correspond, respectively, to the "non-uniform" and "new uniform" approaches to copy raising that we discuss in the next section.)

Note also that the label "direct evidentiality" is somewhat misleading, since 'seem'-reports are arguably only appropriate when the speaker has somewhat *indirect* evidence for the embedded claim. With rain falling around you, you don't say that it *seems* like it's raining, but that it *is*. The term "direct," as used by Rett & Hyams (2014) in this context, refers to the direct acquaintance with subject, which is the source of evidence, not to the directness of the evidence with respect to the embedded claim. See also Asudeh & Toivonen (2017), §3.

[8]This style of story makes immediately relevant two types of cases: first, SLPs that are such that evidence for their application doesn't seem especially tied to an individual's appearance; and second, ILPs that are such that evidence for their application *does* seem closely tied to an individual's appearance. Examples of the former may be stage-level predicates like 'missing' or 'absent' — and indeed, Doran (2015) gives an example with 'missing' (with the appearance verb 'sound') to argue against a uniform analysis (as defined in below in §6) for all CR reports (p. 11). Examples of the latter would be individual-level predicates like 'tall' or 'brunette'. More work is needed to determine to what extent the acquaintance behavior of CR reports tracks a predicate's status as SLP or ILP itself, as opposed to tracking something else that just tends to be correlated with these categories of predicates.

[9]Simplified from Asudeh & Toivonen (2012); see pp. 331, 344, 361.

(8) **Context** A and B walk into Tom's kitchen. There's no sign of Tom, but there are various things bubbling away on the stove and there are several ingredients on the counter, apparently waiting to be used.

a. # Tom seems like he's cooking.

b. $\exists s[seem(s, cooking(tom)) \wedge psource(s) = tom]$

c. It seems like Tom is cooking.

d. $\exists s \exists x[seem(s, cooking(tom)) \wedge psource(s) = x]$

Asudeh & Toivonen (2012), as well as Rett & Hyams (2014), hold that all CR reports are unacceptable in contexts where the matrix subject is not perceived (that is, they all have the specific acquaintance inference), and hence offer what we will call a *uniform analysis*, on which all CR reports assign the p-source role to the matrix subject. They thus take all CR reports to have meanings like the one in (8b).

This uniform p-source analysis is empirically inadequate, as we have seen from the examples of copy raising reports embedding ILPs. And indeed, cases of CR reports that seem acceptable in contexts where the matrix subject isn't perceived (that don't have the specific acquaintance inference) have led Landau (2011) and Doran (2015) to offer a *non-uniform analysis*, on which CR subjects sometimes do and sometimes don't get assigned the p-source role. They thus take some CR reports to be analyzed as in (8b), while others are analyzed just like ES reports, as in (8d).

However, the key question of why some copy raising report is interpreted one way or the other goes unanswered in this literature. It is here that the work pursued in the present paper comes in. Importantly, though, the specific with stage-level generalization defended here only has implications for the semantics of 'seem'-reports under an assumption implicitly endorsed by all of the above-mentioned theorists: namely, that the acquaintance behavior of the reports in question should get encoded in the semantics.

> **Acquaintance/p-source link** The matrix subject of a copy raising 'seem'-report is interpreted as the perceptual source if and only if the report gives rise to the specific acquaintance inference.

Assuming this link, the specific with stage-level generalization lends support to a non-uniform analysis of CR reports. This is because the generalization holds that some CR reports give rise to the specific acquaintance inference while others don't, and hence, given the acquaintance/p-source link, that some CR subjects are interpreted as p-source while others aren't. But more significantly, the generalization offers a systematic answer to the question of which CR reports get interpreted which way — something missing in previous discussion of non-uniform analyses. The p-source interpretation is present, on the current proposal, when its presence allows us to mark a distinction in perceptual evidential situation, and this is the case with 'seem'-reports embedding SLPs, but not those embedding ILPs.

This ends our discussion of the modest upshot of the specific with stage-level generalization. It serves to systematize variable acquaintance behavior across copy raising 'seem'-reports, and thus to underwrite a non-uniform analysis, on which some of these reports specify a p-source, while other don't.

But isolating an assumption is often the first step towards questioning it. We would thus like to conclude this theoretical discussion slightly more radically, by considering whether the acquaintance/p-source link may, in the end, be done without.

Rejecting the acquaintance/p-source link opens up space for a *new uniform analysis*, on which all CR reports have the same underlying semantic structure, while still not making incorrect predictions about the acquaintance inferences associated with the reports. On such a view, the role assigned to the matrix CR subject would be more minimal than the p-source role. To have a label, we'll call this the *evidential source* (e-source) role. The e-source is the source of appearance-based evidence. Crucially, what it takes to be an e-source may vary depending on the claim that the evidence is supposed to be evidence for. Trivially, Tom can be an evidential source for the claim that he's cooking, without being an evidential source one way or the other for the claim that he's an experience cook — say if what he's doing is too basic to show the difference. To capture the results of the specific with stage-level generalization, however, the important possibility is that what it takes for Tom to be an e-source for a stage-level claim about him may be more restricted than what it takes for Tom to be an e-source for an individual-level claim about him. He can be an e-source for the claim that he's cooking only by appearing some way himself; but he can be an e-source for the claim that he's an experienced cook also by being appropriately connected with the scene in the kitchen that suggests as much. More generally, the claim would be that an individual can only be an e-source for its possession of stage-level properties by appearing some way itself; whereas it can be an e-source for its possession of individual-level properties not only by appearing some way itself, but also by being appropriately connected with some other scene that appears some way. So on this approach, there are two very different ways for a 'seem'-report to lack the specific acquaintance inference. Expletive subject reports lack this inference, because nothing in particular is singled out as the e-source. Copy raising reports with embedded ILPs lack this inference too, but not because no e-source is specified. Rather, they lack the inference, because something can be an e-source without being perceived.

The e-source approach captures the variable acquaintance behavior of copy raising 'seem'-reports, without positing different semantic structures for surface-alike sentences. The explanation rests on independent facts about stage- and individual-level properties, and what it takes to be sources of evidence for their presence. However, it faces some challenges. We close by mentioning two.

First, the e-source approach owes more of an explanation for the different evidential behavior of SLPs and ILPs. The claim is that Tom himself can't be the e-source for the claim that he's cooking, in the absent cook case. Still, in that scenario, something — say, the scene in the kitchen — must be an e-source. This is what allows the expletive subject report to be appropriate. So far, so good. But the e-source approach must say something very different about the individual-level example (the "absent experienced cook case," if you will). In that case, Tom — Tom himself, and not just the scene in the kitchen! — can be the e-source for the claim that he's an experienced cook, even if he isn't present. This is what is supposed to explain why the CR report embedding 'an experienced cook' doesn't give rise to the specific acquaintance inference. But why would this be? We would like to have

something more to say here.

Second, the e-source approach faces a challenge from 'seem'-reports without embedded pronouns that co-refer with the subject.

(9) Pavi seems like the baby didn't sleep last night.

There has been some controversy over the acceptability of such *copy free* (CF) reports.[10] We hold, following Landau (2011) and Kim (2014), that they can be acceptable, given the right context. CF reports all give rise to the specific acquaintance inference. For (9) to be appropriate, the speaker must have perceived Pavi, perhaps having walked in looking sleep-deprived, or otherwise showing signs of having been kept up by a baby. Reports like these pose a challenge for the e-source analysis because there are minimal pairs of CR and CF reports, with equivalent embedded clauses, but which differ in their acquaintance requirements.

(10) **Context** A and B walk by their neighbor, Charlie's house. A notices that a ground-floor window has been smashed. The address of the house is 12 Front Street.

a. **A:** Charlie seems like he's been burgled.

b. # **A:** Charlie seems like 12 Front Street has been burgled.

The CF report in (10b) is clearly unacceptable in the context described, where Charlie himself is not perceived by the speaker. (The report could be acceptable in a context where Charlie comes in looking distraught, and the speaker knows that Charlie has two houses, but cares much more about the house at 12 Front Street.) This is unexpected on the e-source approach. To the extent that the CR report in (10a) is acceptable, the e-source analysis holds that this is because Charlie can be an evidential source for the claim that he's been burgled, even if he's not present. But if he's the e-source for that claim, surely he can equally be the e-source for the (known-to-be) extensionally equivalent claim that 12 Front Street has been burgled. The e-source analysis seems to be at a loss to explain why CR and CF reports should have such different acquaintance requirements.[11]

7 Conclusion

The acquaintance inferences associated with 'seem'-reports are variable in a way that has no analogue in the more widely-discussed cases of assertions of about personal tastes and aesthetic matters. Reports with 'seem' call for a distinction between the specific acquaintance inference (that the speaker is acquainted with a specific individual), and the general acquaintance inference (that the speaker is acquainted with something or other relevant to the truth of the embedded claim). Which type of acquaintance inference is associated with a given 'seem'-report depends on the previously unnoticed factor of whether the embedded clause contains a stage-level or

[10] See, e.g., Asudeh & Toivonen (2012), p. 330.

[11] There is the option of holding that while CR 'seem'-reports have e-source matrix subjects, CF reports have proper p-source subjects. Without more explanation, though, this would be a very ad hoc move.

individual-level predicate. We have demonstrated this dependence — the specific with stage-level generalization — with new experimental results, and explained why it makes sense functionally, given the evidential role of 'seem'-reports.

We then discussed the relevance of these results for the semantics of these reports. More modestly, this work fills out a non-uniform analysis of copy raising reports, by giving a systematic way of telling which reports give rise to which kind of acquaintance inference, and hence which reports call for a perceptual source interpretation of the matrix subject, and which don't. More radically, it may suggest rejecting the assumption that the specific acquaintance inference calls for a designated p-source interpretation at all. This latter approach involves giving a semantics of copy raising reports that is general enough to allow for both specific and general acquaintance inferences, depending on the case. We spelled out a version of this approach, making use of an evidential source role to replace the perceptual source one. It faces unresolved challenges, however. Given this situation, we officially rest on the more modest conclusion. Still, we hope exploration of the new uniform e-source approach is instructive in helping map out the space of possible approaches to 'seem'-reports and their acquaintance inferences.

References

Asudeh, A., & I. Toivonen. 2012. Copy raising and perception. *Natural Language & Linguistic Theory* 30.321–380.

Asudeh, A., & I. Toivonen. 2017. A modular approach to evidentiality. In *Proceedings of the LFG '17 Conference*, ed. by M. Butt & T. H. King, 45–65. CSLI Publications.

Carlson, G. N. 1977. *Reference to Kinds in English*. University of Massachusetts, Amherst dissertation.

Doran, D. 2015. A semantic analysis of english copy raising constructions. Master's thesis, McMaster University.

Kim, J.-B. 2014. English copy raising constructions: Argument realization and characterization condition. *Linguistics* 54(1).167–203.

Landau, I. 2011. Predication vs. aboutness in copy raising. *Natural Language & Linguistic Theory* 29.779–813.

Ninan, D. 2014. Taste predicates and the acquaintance inference. In *Proc. Semantics and Linguistic Theory 24*, 290–309.

Pearson, H. 2013. A judge-free semantics for predicates of personal taste. *Journal of Semantics* 30(1).103–154.

Rett, J., & N. Hyams. 2014. The acquisition of syntactically encoded evidentiality. *Language Acquisition* 21(2).173–198.

Rogers, A. 1972. Another look at flip perception verbs. In *Papers from the 8th regional meeting of the Chicago Linguistics Society*, ed. by P. Peranteau, J. Levi, & G. Phares, 303–315.

Rudolph, R. E. 2019. A closer look at the perceptual source in copy raising constructions. In *Proceedings of Sinn und Bedeutung 23*. Forthcoming.

Stump, G. 1985. *The Semantic Variability of Absolute Constructions*. Dordrecht: Reidel.