Generality, genericity and subjective predicates: What propositional attitude verbs, alien viruses, and COVID can tell us[*]

Elsi Kaiser & Haley Hsu
University of Southern California

Abstract: In uttering a subjective opinion like *Donuts are tasty*, is a speaker expressing her own opinion or also making a generalization about other people? While researchers largely agree that generic readings of subjective predicates exist, there is no consensus on how central genericity is for theories of subjective meaning. We report a psycholinguistic study that tests what influences the level of prevalence that comprehenders attribute to opinions, expressed with subjective predicates, about unfamiliar information. Specifically, if you overhear an alien expressing an opinion about an unfamiliar virus (e.g. *The zorgavirus is dangerous*), how many other aliens do you think share this alien's opinion? We find that the perceived generalizability of subjective predicates is modulated by the presence/absence of embedding under propositional attitude verbs (whether the speaker is explicitly mentioned with *I think/consider*) and by participants' extra-linguistic attitudes, namely their anxiety levels about COVID. This work uncovers a new link between subjective predicates and humans' egocentric cognitive biases.

Keywords: subjective predicates, predicates of personal taste, egocentricity, experimental semantics, genericity, perspective-taking, attitude verbs, COVID-19

1 Introduction

Subjective predicates (including predicates of personal taste, PPTs, e.g. *fun, tasty, scary*, (1)) convey opinions, unlike objective predicates (e.g. *organic, plastic*).

(1)   Sam: Donuts are tasty.

However, utterances like (1) are potentially ambiguous in terms of whose opinion they express. It is well-known that subjective predicates favor speaker-oriented interpretations, but under certain circumstances they can also receive a non-speaker-anchored interpretation (e.g. Lasersohn 2005; Anand & Korotkova 2021

i.a. on autocentric vs. exocentric interpretations; Kaiser 2021, 2022 for psycholinguistic experiments). In addition, even when a subjective predicate expresses a judgment that reflects the speaker's opinion, there is debate about whether a generic interpretation is also available (see e.g. Lasersohn 2005; Stephenson 2007; Snyder 2013; Keshet 2020; Pearson 2022).

1.1 Generic interpretations of predicates of personal taste

When Sam says (1), is she expressing her own opinion or making a generalization about people-in-general? It has been suggested that in addition to expressing the speaker's opinion, these utterances can be interpreted generically. In the words of Pearson (2013), such utterances "are used to make statements about whether something is tasty to *people in general*, based on first person experience" (Pearson 2013:121, emphasis added). Relatedly, Moltmann (2010) describes them as expressing a generalization that is "to be true with respect to anyone of the relevant sort (including the interlocutors)" (Moltmann 2010: 210). Bhatt & Izvorski (1997) also discuss generic interpretations of evaluative predicates in the context of implicit arguments in control structures.

However, there is not yet a consensus on the status of generic interpretations. Some researchers appear to assume that a generic reading is always available, while others disagree, and yet others, including Keshet (2020), seem to treat generic and 'speaker-specific' construals as dissociable. Let's take a closer look at this third option first. Consider (2), adapted from Lasersohn (2005) by Keshet (2020: 73).

Focusing just on the exchange in (2a) and (2b), Lasersohn notes that on the one hand, it feels like John and Mary are disagreeing, but on the other hand, it feels that both are in the right,[1] as both are making statements that are true when relativized to their experiences (which Lasersohn formalizes with a 'judge' variable).

(2)   a. John: The chili is tasty.
       b. Mary: No, the chili is not tasty.
       c. John: OK, the chili is not tasty for you (Mary), but it is tasty for me

Keshet instead argues for a distinction between 'tasty for the speaker' and 'tasty in general.' In terms of Lasersohn's notion of a judge variable: the judge variable can be set to the speaker or set to a generic value. According to Keshet, John's initial utterance in (2a) is in principle ambiguous. It could be a speaker-specific use, where John is expressing his own opinion of the chili (judge=speaker). Or it could be a generic use, where John is saying that the chili is 'tasty generically,' tasty to a generic eater (judge=generic).

---

[1] This phenomenon is known as faultless disagreement (Kölbel 2003).

What about Mary's response in (2b)? If the judge in (2a) and (2b) is set to the speaker of the sentence, there is no disagreement, as John and Mary are both simply stating their own experiences (faultless disagreement). But if the judge in (2a) and (2b) is set to be a generic value, then there can be a disagreement: now, John and Mary disagree about what is tasty to a generic eater. In other words, for Mary's response in (2b) to 'count as' a true disagreement, she must be interpreting John's statement as generic, and responding to it with another generic statement.

Then, in (2c), John essentially clarifies that he meant (2a) to convey a speaker-specific reading and says to Mary that "I presume that you based your statement that the chili is not generically tasty on the evidence that the chili is not tasty for you. However, the chili is tasty for me, so perhaps your reasoning is faulty. You leapt to the conclusion that the chili is not generically tasty because it was not tasty to you" (slightly reworded from Keshet 2020: 73).

Keshet attributes the specific-generic ambiguity to a hidden pronoun, rather than a variable. He posits that PPTs "make use of a hidden pronoun (…) [whose] value (…) can be either generic or specific. When specific, it is usually the speaker, but can take other values" (Keshet 2020: 74).[2]

This approach treats the specific and generic readings as distinct and dissociable. A related analysis is proposed by Snyder (2013), who argues that PPTs, by virtue of being individual-level predicates, allow generic construals, and by virtue of being anaphoric implicit-argument-taking predicates (similar to *local*, *apply*), can also receive context-dependent anaphoric interpretations. Relatedly, Stephenson (2007) concludes that "generic readings of predicates of personal taste are generally available, in addition to the normal, judge-dependent reading" (Stephenson 2007: 57). In recent work, Pearson (2022) argues that the covert experiencers of PPTs in predicative contexts like (1) and (2a) obligatorily receive generic readings. In her analysis, this stems from covert pronominal arguments of individual-level predicates needing to be bound by the generic operator.

However, Lasersohn (2005) uses examples like (3a,b) to argue that a generic reading of PPTs is not the only available one.

(3)　　a.　　A paper clip collector: "This is fun, but most people would hate it."
　　　　b. ?? This is not fun at all, although I'm having fun doing it.

Lasersohn (2005) points out that (3a) is felicitous, even though most people (or 'people in general') do not find collecting paper clips to be fun. He notes that the felicity of (3a) seems unexpected if the PPT can only have a generic interpretation.

---

[2] Our study does not hinge on whether subjective predicates contain an implicit pronoun or variable. Regardless of how one analyses subjective predicates (i.e., even if one does not assume a covert element), our results have consequences for the semantics and pragmatics of subjective predicates.

Furthermore, he notes that (3b) sounds odd, which again seems unexpected if the PPT can only have a generic reading. However, Keshet (2020) argues against these concerns by pointing out that if we allow for a hidden pronoun that can receive a specific *or* a generic reading, the correct judgments can be derived.

Furthermore, as noted by Carlson (1977) and also discussed by Snyder (2013), a generic sentence can be acceptable even when the property does not hold of the majority. Consider, for example, 'Ducks lay eggs' or 'Mosquitoes carry malaria'. These are typically judged to be good generic sentences, although only about half of all ducks lay eggs (female ducks) and most mosquitoes do not in fact carry malaria (see e.g. Carlson 1977, Krifka et al. 1995, Leslie et al. 2011, Tessler & Goodman 2019). It's worth noting, though, that these non-majority readings are only possible in certain contexts; 'Ducks are female' sounds odd. Given these properties of generics, Snyder (2013) argues against Lasersohn's objections to a generic approach to PPTs.

In sum, while researchers largely agree that generic readings exist, there is a lack of consensus about how central they are to our understanding of PPTs and how to incorporate generic readings into semantic analyses of PPTs.

1.2 Aims of this paper: Perceived generalizability of subjective predicates

Prior work suggests that predicates of personal taste can (and according to some, must) receive generic readings, which can be viewed as arising from generic quantification over potential judges or implicit arguments, depending on one's analysis. But this does not mean that the majority of people share the speaker's opinion: it is felicitous for a speaker to say *Oysters are delicious* in a context where not everyone, or even the majority, shares this opinion.[3] Ultimately, subjective predicates can be felicitously used in various contexts, ranging from situations where they reflect the opinion of a minority (*Collecting paper clips is fun*) to situations where they reflect a more widely-held opinion (*Listening to music is fun*).

This brings up the question of how comprehenders interpret sentences with PPTs, and subjective predicates more generally – in particular, what influences how widespread comprehenders perceive these opinions to be? Imagine that someone

---

[3] To reflect this, in her (2013) genericity-based account of PPTs, Pearson builds on Moltmann (2010a) and incorporates an *identify with* relation into her analysis. She notes that "To say that *the cake is tasty* is to say that the cake is tasty to every (contextually restricted) individual with whom I identify. The *identify with* relation is intended to model a notion of empathy and is therefore reflexive: I always empathize with myself. In normal conversation, we empathize with our interlocutors, but we may also empathize with other individuals (…)" (Pearson 2013: 123). Thus, the scope of the generalization is contextually restricted and restricted to the individuals that the speaker 'identifies with' (see also Moltmann 2010a,b).

travels to a faraway place and overhears a local say *Churbits are tasty*. The tourist has never heard of churbits before. The tourist can (normally) assume that the speaker finds churbits tasty, but can they draw any inferences about how many other people share this opinion (how widely-held the opinion is)? Or what if the speaker had said *I think that churbits are tasty*? Does embedding under an attitude verb influence the inferences about how widely-held the opinion is?

These are the kinds of questions that we investigate in the experiment reported in this paper, which builds on some of our earlier work (Kaiser & Lee 2019). More specifically, we test whether and how linguistic factors (matrix clauses vs. embedding under the propositional attitude verbs *think* and *consider*), as well as non-linguistic egocentric biases, guide the level of generalizability that comprehenders attribute to subjective predicates.

In the work reported here, we use the notion of *how widely-held an opinion* is (i.e. *how generalizable it is*) as an admittedly rough 'proxy' for the generic interpretation of subjective predicates. In other words, we make the simplifying assumption[4] that if participants perceive an opinion as highly generalizable, i.e., if they think that a lot of people share that opinion, then this can indicate that a generic interpretation of the subjective predicate expressing that opinion is available. Put differently, if a comprehender, in a particular context, can easily interpret a subjective predicate generically (e.g. if they interpret *tasty* as *tasty to people in general* or *tasty for a generic taster*), then we should find that the comprehender will regard the opinion as being widely-held (i.e., held by people in general).

It is also important to note that our focus is purely on the *comprehender*. In other words, if you hear or read an utterance with a subjective predicate, what kinds of inferences (if any) do you as the comprehender make about the prevalence of the opinion expressed by the predicate? This is not the same as investigating what the *speaker* thinks about how widespread their opinion is. The study reported here does not seek to test how widespread speakers think their opinions are, or whether and how speakers formulate their utterances to reflect their perception of how many other people would agree with them. We leave these questions for future work.

The structure of the paper is as follows. In Section 2, we review prior work on effects of embedding subjective predicates under attitude verbs and introduce the Opinion-Holder Explicitness Hypothesis. According to this hypothesis, opinions expressed with subjective predicates in matrix contexts are interpreted as more widely held, more generalizable, than opinions expressed with subjective predicates embedded under *I think* or *I consider* (i.e., in contexts where the first-person opinion-holder is linguistically overt). In Section 3, we review cognitive psychology work on humans' egocentricity bias which shows that we often project

---

[4] Treating perceived high generalizability/prevalence as an indication of the availability of a generic interpretation is clearly an over-simplification. However, we also note that other work points to a deep connection between genericity and prevalence/likelihood (e.g. Tessler & Goodman 2019).

our own attitudes and opinions to others and overestimate the extent to which our beliefs, attitudes and knowledge are shared by others. We also introduce the Shared Opinion Hypothesis. According to this hypothesis, when a participant agrees with an opinion, they will view this opinion as more widely-held by people in general (perceive it as more generalizable) than if they do not agree with the opinion (in which case they will think the opinion is less common, less generalizable). This prediction is based on the idea that people's egocentric biases guide how they interpret opinions matching vs. mismatching their own opinions.

In Section 4, we test these hypotheses experimentally. Our results are reported in Section 5, and a brief discussion is in Section 6. We find support for both the Opinion-Holder Explicitness Hypothesis and the Shared Opinion Hypothesis, and provide (to our knowledge) the first systematic investigation of how non-linguistic individual differences guide humans' interpretation of subjective predicates.

2 Embedding under propositional attitude verbs

The first aim of the experiment reported here is to test whether explicitly anchoring an opinion to an opinion-holder by means of a propositional attitude verb influences how others interpret the opinion. In addition to being expressed in matrix clauses (4a), subjective predicates can be embedded under attitude verbs like *think* (4b). Embedding a subjective predicate under *think* or *consider* spells out whose opinion the adjective describes (e.g. Stephenson 2007; Nouwen 2007; Patel-Grosz 2015). In saying (4b), the speaker explicitly self-identifies as the opinion-holder.

(4)   a. The chili is tasty.
        b. I think the chili is tasty.

In (4a), the speaker's identity as the opinion-holder is not explicit, but there is nonetheless an implicit assumption that the speaker is the one who holds this opinion (e.g. Lasersohn 2005; Stephenson 2007; Ninan 2014). For example, if I say *the cake is tasty*, the default assumption is that I have actually tasted the cake and liked it (e.g. Pearson 2013).

If matrix clauses like (4a), with no explicit opinion-holder, can convey not only the speaker's specific opinion but also a generalization about others, one might predict that hearing (4b), with an explicit opinion-holder, will lead a comprehender to interpret the opinion that the chili is tasty as being held by less people (less generalizable) than if the comprehender had heard or read (4a).[5] We refer to this as

[5] The predicted difference between (4a-b) can be derived from Gricean principles (Grice 1975). For example, Horn (1984)'s Division of Pragmatic Labor predicts that a longer expression licenses an inference that the unmarked situation does not hold: If I say *I think* or *I consider* (instead of using a shorter form), this licenses an inference that not everyone shares my opinion. In other words, use of

the **Opinion-Holder Explicitness Hypothesis**.

In our earlier experimental work (Kaiser & Lee 2019), we provide initial support for this hypothesis. We showed that explicit mention of the speaker by means of *I think* results in comprehenders perceiving opinions as being less generalizable (less widely held), compared to when the same opinions are expressed as plain matrix sentences. Putting it somewhat differently, that earlier study suggests that explicit mention of the speaker boosts the specific (speaker-anchored) reading, while plain matrix sentences are more easily construed as having broader referential scope (generic reading).

In the current study, we test if this result replicates in a novel context and with a large set of linguistically naïve participants, and whether it extends to another attitude verb, *consider*, that differs semantically from *think* (see e.g. Lasersohn 2007; Moltmann 2010b; Kennedy & Willer 2022 on *consider*). Although the Opinion-Holder Explicitness Hypothesis, in its present form, makes no predictions about *think* vs. *consider*, one might expect that embedding a subjective predicate under *consider* might make it less generalizable than embedding under *think*.

3 Consequences of humans' egocentric biases?

As mentioned above, one of the aims of the present work is to assess whether a linguistic factor – explicit mention of the opinion-holder – modulates the perceived generalizability of opinions. The second aim of this work is to test whether nonlinguistic individual differences (participants' own attitudes) can modulate the perceived generalizability of opinions, given that humans tend to over-generalize their own attitudes to others (e.g. Epley & Caruso 2009).

Although we are able to think about other people's mental states and realize that they are distinct from our own (Theory of Mind), research in cognitive psychology shows that people nevertheless tend to be egocentrically biased. Possessing the ability to take on someone else's perspective does not mean that people consistently do so (see e.g. Epley & Caruso 2009 for an overview).

Although a large body of work shows that real-time language processing can be guided by perspectival information very early on (see Brown-Schmidt & Heller 2018 for an overview), non-linguistic research indicates that we often fail to fully use our perspective-taking abilities and have an egocentric bias. For example, humans often overestimate the extent to which their beliefs, attitudes and knowledge are also shared by others (e.g. Birch & Bloom 2007; Keysar & Barr 2002; Krueger & Clement 1994). Some early work on egocentric bias found that

---

a marked expression licenses a more marked meaning. Levinson (2000)'s M principle essentially derives the same effect.

when participants felt angry or anxious/afraid, they judged photographs and stick figures as more hostile or anxious, compared to control groups who did not feel angry or anxious/afraid (e.g. Feshbach & Feshbach 1963; Feshbach et al. 1963). This suggests that people generalize from their own emotional states to others, and assume that others experience those states as well.

Projection from oneself to others has also been observed for biological drive states such as thirst. For example, when participants who were thirsty (due to vigorous exercise) were asked to predict how much a group of hungry and thirsty hikers would be bothered by hunger and thirst, the thirsty participants thought that the hikers would be more bothered by thirst than did other participants who were not thirsty (Van Boven & Loewenstein 2003). In the words of Epley & Caruso (2009), "Such egocentric projection of drive states suggests again that people make predictions about how others will feel by first imagining what they themselves would feel in the others' situation and do not adjust sufficiently to overcome their own egocentric perspective" (Epley & Caruson 2009: 303).

Thus, in various domains, we over-generalize our own experiences and opinions to others. To quote Snyder, Stephan & Rosenfield (1978: 92), "If one believes that one's own response is the natural reaction, then others who are confronted with the same reality should respond in the same way – unless they are somehow deviant. Consequently, one may believe that others share one's opinions more than they really do." Humans' over-reliance on the self when thinking about others' opinions, characteristics, feelings, decisions and behavior has been extensively researched. This broad phenomenon is referred to as *egocentric attribution* (Heider, 1958), *(social) projection* (e.g. Katz & Allport 1931; Holmes 1968), *egocentric bias* (Epley et al. 2004), *assumed similarity* (Cronbach 1955), *false consensus* (e.g. Ross, Green & House 1977) and *looking-glass perception* (e.g., Fields & Schumann 1976).

Let us now return to language, and to subjective predicates. In light of the large body of work on egocentric over-generalization showing that our own mental states can 'leak into' our representations of others' mental states, we predict that the level of generalizability that a particular comprehender attributes to a subjective predicate can be influenced by that comprehender's own mental state. In very simplistic (and categorical) terms, one might predict that whether Person X interprets Person Y's statement that "The chili is tasty" specifically, as reflecting only Person Y's own opinion, or generically, as expressing a claim that people-in-general find the chili tasty, could depend on person X's level of hunger.

Building on this, we assume a more nuanced approach and propose the **Shared Opinion Hypothesis**. According to this hypothesis, when a comprehender agrees with the opinion being expressed, they will view the opinion as being more widely-held by people in general (perceive it as more generalizable) than if they do not agree with the opinion (in which case they will think that the opinion is less

common, perceive it as less generalizable). This prediction stems from the comprehender's bias for egocentric projection guiding how they interpret opinions matching vs. mismatching their own views.

Let's consider example (5). According to the Shared Opinion Hypothesis, if Sam says 'These oysters are tasty' and Alex herself loves oysters, Alex is more likely to interpret Sam's opinion as being prevalent, widely-held by people in general. This is because Alex is likely to overproject/overgeneralize her own liking of oysters to other people (as prior work shows, humans engage in egocentric projection), and thus when she encounters an opinion like (5), the egocentric projection ('others like oysters too') will favor her interpreting Sam's opinion as being widely-held (i.e., she will perceive it as highly generalizable).

(5)   Sam: These oysters are tasty.
a.   Alex loves oysters and overgeneralizes this positive view of oysters to
        others => Alex is likely to interpret Sam's opinion as being widely-held
b.   Alex hates oysters and overgeneralizes this negative view of oysters to
        others => Alex is likely to interpret Sam's opinion as being less widely-held

On the other hand, if Alex hates oysters, we predict that she is likely to interpret Sam's opinion as being relatively *less* widely-held, *less* generalizable. This is because, due to egocentric projection, Alex assumes that most people are like her and dislike oysters. Thus, she is more likely to interpret Sam's claim that the oysters are tasty as being a more 'narrow' opinion, held only by a small number of people.

This novel prediction is rooted in existing nonlinguistic work on human cognition and perspective-taking but has not previously been tested with subjective predicates. Note that this approach does not make any claims or predictions about Sam's intended meaning, wording choices or expectations about how others would interpret her statement. We leave these questions for future work.

So far, we have been considering a situation where both Sam and Alex have direct experience with and specific opinions about oysters. In such a case, there is a clear connection: Sam made a statement specifically about oysters and Alex has opinions about oysters. However, what if Sam says that the clams – a kind of edible mollusk similar to oysters yet different – are tasty, but Alex has never seen or tasted clams, has no opinions about how clams taste, and only has experience with oysters? Would we still find effects of Alex's attitudes about oysters influencing the level of generalizability that she attributes to Sam's opinion about the clams?

In other words, are these effects powerful enough to arise even in situations where the speaker expresses an opinion about one thing and the comprehender has experience with a *related-but-different* thing (and is thus egocentrically projecting her opinion about that related thing)? The study that we report in this paper explores this second kind of situation, in order to shed light on how far-reaching and how robust potential Shared Opinion effects are. In particular, we tested what level of

generalizability participants attribute to opinions about unfamiliar viruses, to see whether this is influenced by participants' own attitudes about a related-but-distinct topic, namely the coronavirus/COVID-19.

4 Experiment

We conducted a psycholinguistic study to test how linguistic form (specifically, the presence/absence of embedding under an attitude verb) and non-linguistic individual differences (specifically, individuals' anxiety levels about COVID) influence their perception of opinions about other (unfamiliar) viruses – in particular, how generalizable such opinions are. When you hear someone express an opinion about a new virus, how widely-held do you think that opinion is? In our study, we used an alien context, with aliens on an alien planet expressing opinions about viruses unfamiliar to our participants, to prevent participants from using real-world knowledge about specific viruses. Our use of an alien context follows the approach taken by Kaiser & Lee (2019). By using unfamiliar viruses and an alien context, we ensure that all participants are equally unfamiliar with the alien world. This allows us to gain a clearer view of the phenomena we are investigating, by removing potential variation stemming from other sources.

To measure the perceived generalizability of opinions – i.e. how widespread participants think an opinion is – participants typed in a number indicating how many other aliens (out of 100) they think share the alien speaker's opinion. Thus, this task essentially asks: To what extent do *our participants* think the opinion is held by aliens-in-general – i.e., how available is the generic reading? Note that we are *not* asking participants to assess how widespread *the alien speaker* thinks an opinion is (and we are also not looking at how widespread an opinion actually is, since we are testing an imaginary context). Rather, our focus is on the comprehension side, and we ask: When participants encounter an opinion about an unfamiliar virus, how does the presence/absence of an embedding attitude verb and the comprehender's own egocentric attitudes about a related topic (COVID) influence how widespread the participant thinks the opinion is.

4.1 Participants

Participation was done remotely over the internet, and participants were recruited via Prolific. The median time to complete the experiment was about 30 minutes. 234 U.S.-born native English speakers participated. Two participants were excluded for self-reporting a vision impairment (as this could impact their access to information about the pandemic). Following recommendations by Chmielewski & Kucker (2020) for web-based experiments, we used a multi-layered approach to detect inattentive participants. Thus, 72 participants were excluded for failing the

required threshold for attention checks that asked people what question they had seen on the preceding screen. We used these attention checks because, for the data to be interpretable, participants need to pay attention to the questions being asked on target and filler trials, and thus (following prior work in our lab), we excluded all participants who made two or more mistakes on ten multiple-choice questions asking about what question had been shown on the immediately preceding screen. We used this kind of task and a fairly strict threshold for two main reasons: First, the critical question being asked on target trials (see Fig. 1) does not fit question types frequently used in online experiments (i.e., the task is presumably not familiar to participants) and thus it is important that participants pay attention to our task rather than relying on or analogizing from what they did in earlier experiments. Second, the dependent variable (a number from 0 to 100) does not easily allow us to detect participants who might be responding randomly or misunderstood the task. Thus, we rely on the attention checks to ensure that participants were paying attention to the questions they were being asked (i.e., doing the intended task).

In addition, 21 further participants were excluded for failing to answer correctly five practice questions designed to tap into whether people understand the task, and another 8 participants were excluded for failing unambiguous attention checks (lower than 70% accuracy) about nonce words used on filler trials (e.g. asking about how many aliens would agree that the word SINTOR has more letters than the word PLAD). Thus, in what follows, we report data for 131 participants.

4.2 Design and materials

We tested whether perceived generalizability of opinions is influenced by (i) linguistic form (matrix clauses vs. embedded under *think/consider*) and (ii) opinion type: anxious vs. non-anxious (Table 1). Figure 1 shows an example of a target trial. The text provides a reminder of the alien planet context and explains that the critical sentence is something that one alien says to another. The task is to indicate how many of 100 randomly-selected aliens share the opinion expressed by the speaker on each trial (for example, that the zorgavirus is dangerous).

We tested opinions expressed in matrix clauses as well as embedded under *think* and *consider*. We tested both *think* and *consider* because (i) we wanted to assess the Opinion-Holder Explicitness Hypothesis and to see whether Kaiser & Lee's results for embedding under *think* replicate with *consider*, and because (ii) we wanted to test, more speculatively, whether the semantic differences between *think* and *consider* have consequences for opinions' perceived generalizability. In both cases the first-person speaker is explicitly mentioned, but as noted by Kennedy & Willer (2022), *consider* differs from *think* in not allowing fully objective information. Thus, we wanted to see whether *consider* results in comprehenders treating the opinion as less generalizable than the more neutral *think*.

632

We are visiting an alien planet. You overhear one of the aliens say:

"The zorgavirus is dangerous."

If we randomly select 100 aliens from this planet, how many of them do you think **share this alien's opinion** about the zorgavirus?

> [   ]

Figure 1   Screenshot showing an example item

matrix	The zorgavirus is dangerous$_{anxious}$ / harmless$_{non-anxious}$
think	I think that the zorgavirus is dangerous$_{anxious}$ / harmless$_{non-anxious}$
consider	I consider the zorgavirus to be dangerous$_{anxious}$ / harmless$_{non-anxious}$

Table 1   An example target item

In terms of opinion types, we tested 'anxious' and 'non-anxious' opinions. Anxious opinions express fear or worry about a virus or its health consequences as well as opinions regarding the importance of projecting oneself from (serious) infection by means of vaccines or antiviral medication. (They might also be called 'careful' opinions.) Non-anxious opinions are the opposite, in that they downplay risks of the virus, dismiss the need for health precautions and criticize medical treatments such as vaccines and antivirals. (These opinions could also be called 'dismissive'.) We focus on the anxious/non-anxious distinction, because we wanted to see whether participants' levels of anxiety about COVID would influence how they perceived the generalizability of anxious vs. non-anxious opinions.[6]

The study included 30 targets. We used the predicative syntactic frame shown in Table 1 (e.g. *X is / to be ADJ*). Each of the 30 targets used a different, made-up

---

[6] The study was designed so that participants encountered anxious and non-anxious opinions expressed with a mix of negative and positive adjectives (e.g. anxious / careful opinion with a positive adjective: *The antiviral medication for the kistaaciviru is fantastic*; anxious / careful opinion with a negative adjective: *The voggare virus is vicious*; non-anxious / dismissive opinion with a positive adjective: *The voggare virus is mild*; non-anxious / dismissive opinion with a negative adjective: *The antiviral medication for the kistaacivirus is useless*). In other words, adjective valence (positive vs. negative) was not confounded with opinion type. Furthermore, adjective valence was contextually defined (e.g. *vicious* is relatively more negative than *mild*, when talking about a virus). Here, we analyze the data in terms of opinion type, as the Shared Opinion Hypothesis makes predictions about opinion type, not about adjective valence.

virus name that was designed to sound plausible/realistic (e.g. *zorgavirus, ligaroxia virus, tomoplissa virus, sompidovirus*). We also tested 30 positive and 30 negative subjective predicates (e.g. *intimidating, frightening, fascinating, irritating, awful, exciting, harmless, dangerous, terrible*). While many of these subjective predicates are predicates of personal taste, some others may be better classified as multidimensional subjective predicates. When used in the predicative frame and in the context of viruses, the predicates we used are most easily interpreted as individual-level, not stage-level (see Pearson 2022 for related discussion).

The study also included 45 fillers, designed to elicit a range of numerical responses. Some filler sentences and questions were objective (e.g. asked about how many aliens would know something), whereas others were more subjective and asked about how many aliens would agree with something, to create a mix of subjective and objective trials. Some contained nonce words. We also included ten attention checks asking about what question had been shown on the immediately preceding screen, as well as six unambiguous attention checks about nonce words (as mentioned in the Participants section).

4.2.1 COVID anxiety scores

After participants completed the main experiment, they answered a series of background questions, including questions probing whether the COVID pandemic made them anxious or afraid. We used a series of questions, adapted from prior work on the pandemic's effects on mental well-being, e.g. Conway et al. (2020), to compute an anxiety score for each person. The maximum possible anxiety score is 50. Participants read 10 statements and rated them on a five-point scale (1 = strongly disagree, 5 = strongly agree). As is commonly done, some statements were designed to be reverse-scored (e.g. (d,e)). The ratings for the sentences were summed so that a higher score indicates a higher COVID-related anxiety level.

(6)   Examples from COVID anxiety section (see Conway et al. 2020)
       a. Thinking about COVID-19 makes me anxious.
       b. The COVID-19 pandemic impacted my mental health negatively.
       d. I am NOT worried about COVID-19.
       e. The COVID-19 pandemic has NOT made me feel any better or worse than
          I did before.
       f. I am worried that I or people I love will get sick from COVID-19.

By calculating a COVID anxiety score for each participant, we can test whether individual differences in COVID anxiety levels relate to the level of generalizability that participants attribute to opinions about other viruses, and whether this depends on whether the subjective predicates express an anxious or non-anxious opinion.

4.3 Procedure

Participants completed the study remotely over the internet, using Qualtrics (Provo, UT). They first completed the main part of the study, consisting of the target and filler trials about aliens on an alien planet, and then answered the COVID-related background questions. Data collection took place in mid-to-late August 2022 – in other words, at a point in time where the peak of the pandemic was over in the U.S., but some people were still careful about travel and crowded indoor settings.

5 Results

We first report how the perceived generalizability of subjective predicates is influenced by embedding (vs. lack thereof) under *think* and *consider*, and then turn to how participants' anxiety levels regarding COVID influence their perception of the generalizability of subjective claims about unfamiliar viruses.

5.1 Embedding under attitude verbs

Figure 2 shows the perceived generalizability of subjective opinions about the novel viruses; the y-axis shows how many aliens participants think share the opinion stated by the speaker. Visually, it is clear in Figure 2 that subjective predicates in matrix contexts are rated as expressing more widely-held opinions than subjective predicates embedded under *I think* or *I consider*. Indeed, this pattern is confirmed by linear mixed effects modeling (with the *lme4* package in *R*). When the subjective predicate is embedded under an attitude verb that mentions the speaker (*I think/consider*), the opinion is judged less generalizable than when the subject predicate is in a matrix clause (matrix vs. *think*: t = -3.603, p < 0.001; matrix vs. *consider*: t = -4.404, p < 0.001). Subjective predicates embedded under *think* and *consider* do not differ (t = - 0.802, p > 0.4). Thus, participants think that relatively more aliens share the opinion when the opinion is in a matrix clause, a finding in line with the Opinion-Holder Explicitness Hypothesis. This replicates our earlier findings (Kaiser & Lee 2019), obtained in a context unrelated to viruses or COVID.

Perhaps unexpectedly, in the present study, *think* and *consider* (this second verb was not tested by Kaiser & Lee 2019) pattern alike in terms of the opinion's perceived generalizability. This could be taken as an indication that any kind of explicit mention of the speaker mention boosts its prominence (and thus lowers the perceived generalizability of an opinion), regardless of fine-grained semantic differences between the verbs. In this regard, it further strengthens the reasoning underlying the Opinion-Holder Explicitness Hypothesis.
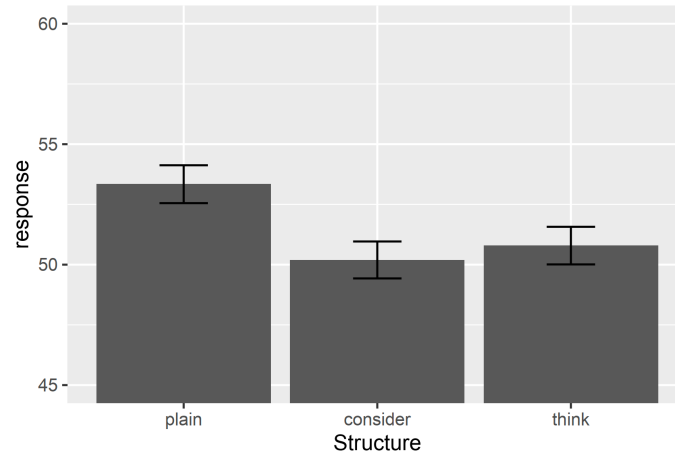
Figure 2 Perceived generalizability. (Y-axis: How many aliens do you think share this alien's opinion?) Error bars show +/- 1 SE.

Because the Opinion-Holder Explicitness Hypothesis makes no predictions about effects of opinion type on generalizability, both opinion types are grouped together in Figure 2. Further analysis shows that for both of the opinion subtypes, the matrix conditions yield numerically higher rates of generalization than the two embedded conditions. In other words, the basic patterns visible in Figure 2 obtain with both anxious and non-anxious opinions.

Furthermore, it is interesting to note that anxious opinions about viruses are rated overall as more generalizable than non-anxious opinions (lmer, t = 49.83, p < 0.0001). In other words, when participants are asked to indicate how many other aliens they think share opinions that reflect a careful, anxious stance towards the virus (e.g. *the zorgavirus is dangerous*), they report a higher number of agreeing aliens than for opinions that reflect a dismissive, non-anxious stance towards the virus (e.g. *the zorgavirus is harmless*). We made no explicit predictions about relative generalizability levels of anxious vs. non-anxious opinions, but this result suggests that participants tend to regard the aliens as having a careful, anxious stance towards viruses. In light of the well-known human tendency for egocentric projection and the unprecedented nature of the COVID pandemic, this is not surprising. Presumably, participants are anxious about COVID (indeed, we will see below that the average COVID anxiety score is relatively high) and thus assume that others will also be anxious about other viruses.

5.2 Individual differences

Now, let us turn to whether people's COVID anxiety scores relate their perception of the generalizability of anxious and non-anxious opinions about other viruses. Overall, the mean COVID anxiety score was 27.8 (range 10-48, median 28, SD 9.9). Figure 3 provides a visualization showing how each participant's anxiety score (on the y-axis; higher scores mean more anxiety) relates to how they perceive the generalizability of anxious and non-anxious opinions about unfamiliar viruses (red circles and blue triangles respectively).[7] Each participant contributes two dots (on the same horizontal y-axis level, corresponding to the participant's COVID anxiety score): one red circle and one blue triangle. The x-axis shows how many aliens participants think share the speaker's opinion.

Non-anxious opinions (blue triangles) cluster to the left of anxious opinions (red circles). We see an asymmetry that reflects what was mentioned earlier: participants view anxious opinions as more widely-held (further to the right) than non-anxious opinions (further to the left). In addition, we see the cluster of red circles and the cluster of blue triangles start out fairly close together near the bottom of the graph (bottom of the y-axis) and then separate from each other more and more as we move upwards and the y-axis values increase. This yields a V-shape that can be seen in Figure 3. This suggests that participants with lower COVID anxiety scores (lower values on the y-axis) do not show as much differentiation between anxious and non-anxious opinions as participants with higher COVID anxiety scores (higher values on the y-axis).

To assess these patterns statistically, we again used linear mixed effects regression (lmer), by fitting a model that now also includes each participant's COVID anxiety score. Before getting into individual differences, we note that this model again confirms that anxious opinions are rated as significantly more generalizable than non-anxious opinions ($t=6.74$, $p< 0.001$). Most relevantly, this effect increases with participants' COVID anxiety level (opinion type x anxiety level interaction, $t=10.86$, $p< 0.001$). This is what yields the general the V-shape.

In essence, the more anxious someone is about COVID, the more likely they are to interpret (i) anxious-opinion sentences as more generalizable and (ii) non-anxious-opinion sentences as less generalizable. The less anxious someone is about COVID, the smaller the difference in how they perceive anxious vs. non-anxious opinions. Thus, as predicted by the Shared Opinion Hypothesis, if someone encounters a worrisome opinion about an unfamiliar virus (e.g. *the zorgavirus is dangerous*) that 'matches' their own opinion on a similar-but-not-identical topic (*COVID-19 is dangerous*), this is associated with the participant perceiving the opinion about the unfamiliar virus as being more widely held.

---

[7] As the Shared Opinion Hypothesis does not make predictions about matrix clauses vs. embedding under *think/consider*, these three conditions are all grouped together in the analyses in this section.
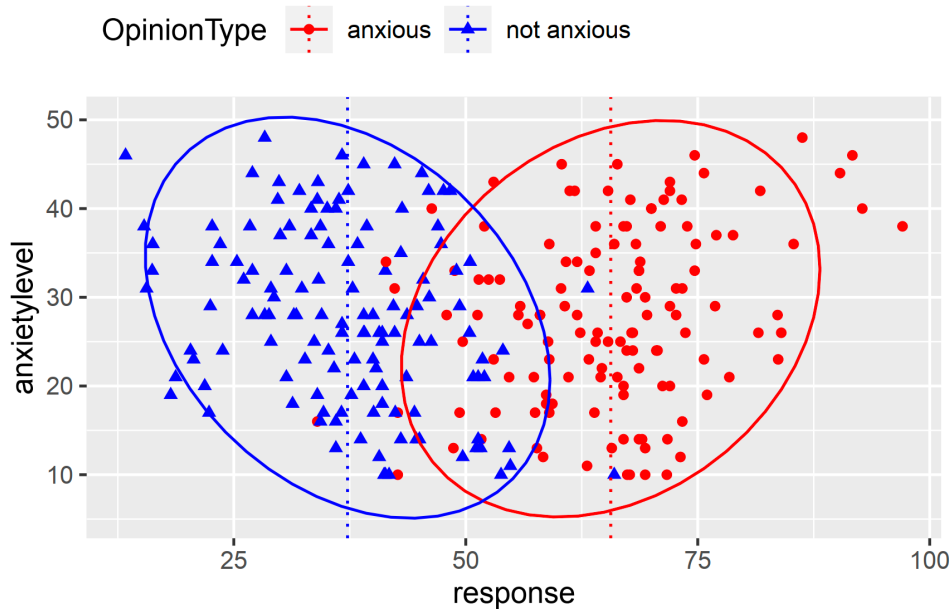
Figure 3 Perceived generalizability (x-axis = how many aliens share the opinion) of sentences conveying anxious/non-anxious opinions, relative to participants' COVID anxiety scores (y-axis). (Dotted vertical lines indicate the mean generalizability of anxious and non-anxious opinions. Ellipses were drawn with *stat_ellipse(type = "t")*)

6 Discussion

Subjective predicates can be used in various contexts, ranging from situations where they reflect the opinion of a relatively small group (e.g. *Blue cheese is tasty*) to situations where they reflect a more widely-held opinion (e.g. *French fries are tasty*, see e.g. Kaiser & Rudin 2020). This brings up the question of how comprehenders interpret sentences with subjective predicates – in particular, in the absence of first-hand experience or specific real-world knowledge, what influences the level of prevalence that comprehenders attribute to opinions? If you hear someone express an opinion about something that you are unfamiliar with, do you make any inferences about how many other people might share that opinion?

   We report a psycholinguistic experiment testing how linguistic factors (matrix clauses vs. embedding under the propositional attitude verbs *think* and *consider*), as well as non-linguistic egocentric biases, guide the level of generalizability that

comprehenders attribute to subjective predicates. More specifically, we investigate two distinct hypotheses. First, according to the *Opinion-Holder Explicitness Hypothesis*, opinions expressed with subjective predicates in matrix contexts are interpreted as more widely held, more generalizable, than opinions expressed with subjective predicates embedded under *I think* or *I consider* (i.e., in contexts where the first-person opinion-holder is linguistically overt). Second, according to the *Shared Opinion Hypothesis*, when a participant agrees with an opinion, they will view this opinion as more widely-held (perceive it as more generalizable) than if they do not agree with the opinion (in which case they will think the opinion is less common, less generalizable). This is based on the idea that people's egocentric bias guides how they interpret opinions matching vs. mismatching their own views.

Using an alien planet context, we tested whether participants' interpretations of opinions about novel viruses are influenced by (a) whether or not those opinions are embedded under attitude verbs and by (b) the participants' own attitudes about another virus, namely COVID-19. If you overhear an alien expressing an opinion about an unfamiliar virus (e.g. The *zorgavirus is dangerous*), how many other aliens do you think share this alien's opinion, and does this depend on how your attitudes about COVID-19?

The results of our experiment support both the Opinion-Holder Explicitness Hypothesis and the Shared Opinion Hypothesis. Participants think that relatively more aliens share the opinion when the opinion is in a matrix clause (compared to when it is embedded under *think/consider*), a finding in line with the Opinion-Holder Explicitness Hypothesis. We also find effects of participants' own attitudes. As predicted by the Shared Opinion Hypothesis, if someone encounters a worrisome opinion about an unfamiliar virus (e.g. *the zorgavirus is dangerous*) that matches their own opinion on a similar-but-not-identical topic (*COVID-19 is dangerous*), this is associated with the participant perceiving the opinion about the unfamiliar virus as being more widely held.

Although many issues remain open for further investigation, this work points to new links between subjective predicates, individual differences, and humans' egocentric cognitive biases, and raises questions about how to capture these effects within current accounts of subjective meaning.

References

Anand, Pranav & Korotkova, Natasha. 2022. How to theorize about subjective language: a lesson from 'de re'. *Linguistics and Philosophy* 45, 619-681. https://doi.org/10.1007/s10988-021-09331-0

Bhatt, Rajesh & Izvorski, Roumyana. 1997. *Genericity, Implicit Arguments, and Control.* Paper presented at the Student Conference in Linguistics (SCIL) VII.

Birch, Susan & Paul Bloom. 2007. The curse of knowledge in reasoning about false beliefs. *Psychological Science* 18(5), 382-386. https://doi.org/10.1111/j.1467-9280.2007.01909.

Brown-Schmidt, Sarah, & Heller, Daphna. 2018. Perspective-taking during conversation. In Gareth Gaskell & Shirley-Anne Rueschemeyer (eds.), *Oxford Handbook of Psycholinguistics*. Oxford: Oxford University Press.

Carlson, Gregory. 1977. Reference to Kinds in English'. PhD, UMass at Amherst.

Chmielewski, Michael & Kucker, Sarah C. 2020. An MTurk Crisis? Shifts in Data Quality and the Impact on Study Results. *Social Psychological and Personality Science* 11(4), 464-473. https://doi.org/10.1177/1948550619875149

Conway, Lucian G., III, Woodard, Shailee R., & Zubrod, Alivia. 2020. *Social Psychological Measurements of COVID-19: Coronavirus Perceived Threat, Government Response, Impacts, and Experiences Questionnaires*. https://psyarxiv.com/z2x9a/

Cronbach, Lee J. 1955. Processes affecting scores on "understanding of others" and "assumed similarity". *Psychological Bulletin* 52(3), 177-193. https://doi.org/10.1037/h0044919

Epley, Nicholas & Caruso, Eugene M., 2009. Perspective taking: Misstepping into others' shoes. In Keith Markman, Klein, William Klein & Julie Suhr (eds.), *Handbook of imagination and mental simulation*, 295-309. New York, NY: Psychology Press.

Epley, Nicholas, Keysar, Boaz, Van Boven, Leaf & Gilovich, Thomas. 2004. Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology 7*(3), 327-339. https://doi.org/10.1037/0022-3514.87.3.327

Feshbach, Seymour, & Feshbach, Norma. (1963). Influence of the stimulus object upon the complimentary and supplementary projection of fear. *The Journal of Abnormal and Social Psychology*, 66(5), 498-502. https://doi.org/10.1037/h0047016

Feshbach, Seymour, Robert D. Singer & Norma Feshbach. 1963. Effects of anger arousal and similarity upon the attribution of hostility to pictorial stimuli. *Journal of Consulting Psychology* 27(3), 248-252. https://doi.org/10.1037/h0048475

Fields, James & Schumann, Howard. 1976. Public Beliefs About the Beliefs of the Public. *Public Opinion Quarterly* 40(4), 427-448. https://doi.org/10.1086/268330

Grice, H. Paul. 1975. Logic and Conversation, in Peter Cole & Jerry Morgan (eds.), *Syntax and Semantics, 3: Speech Acts*, 41-58. New York: Academic Press.

Heider, Fritz. 1958. *The psychology of interpersonal relations*. New York, NY: Wiley.

Holmes, David S. 1968. Dimensions of projection. *Psychological Bulletin,* 69(4), 248-268. https://doi.org/10.1037/h0025725

Horn, Larry R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicatures. In Deborah Schiffrin (ed.), *Meaning, Form, and Use in Context*, 11-42. Washington: Georgetown University Press.

Kaiser, Elsi. 2021. Consequences of sensory modality for perspective-taking: Comparing visual, olfactory and gustatory perception. *Frontiers in Psychology* 12, 701486. https://doi.org/10.3389%2Ffpsyg.2021.701486

Kaiser, Elsi. 2022. Do perspective-sensitive anaphors and subjective adjectives exhibit perspectival uniformity? An experimental investigation. *Glossa* 7(1). https://doi.org/10.16995/glossa.5768

Kaiser, Elsi & Lee, Sarah Hye-yeon. 2019. On the generalizability of subjective opinions: Predicates of personal taste and multidimensionality. Talk given at *Generative Linguistics in the Old World (GLOW) 42,* University of Oslo, Norway, May 10, 2019.

Kaiser, Elsi & Rudin, Deniz. 2020. When faultless disagreement is not so faultless: What widely-held opinions can tell us about subjective adjectives. *Linguistics Society of America (LSA)* 5(1), 698-707. https://doi.org/10.3765/plsa.v5i1.4757

Katz, Daniel & Allport, Floyd H. 1931. *Student attitudes.* Syracuse, NY: Craftsman Press

Kennedy, Chris, Willer, Malte. 2022. Familiarity inferences, subjective attitudes and counterstance contingency: towards a pragmatic theory of subjective meaning. *Linguistics and Philosophy* 45, 1395-1445. https://doi.org/10.1007/s10988-022-09358-x

Keshet, Ezra. 2020. A matter of taste. In Chungmin Lee & Jinho Park (eds.), *Evidentials and modals,* 69-81. Current Research in the Semantics/Pragmatics Interface, Vol. 39. Brill. (manuscript date 2005)

Keysar, Boaz, & Barr, Dale J. 2002. Self-anchoring in conversation: Why language users do not do what they "should." In Thomas Gilovich, Dale Griffin, & Daniel Kahneman (eds.), *Heuristics and biases: The psychology of intuitive judgment*, 150-166. New York: Cambridge University Press.

Kölbel, Max. 2003. Faultless Disagreement. In *Proceedings of the Aristotelian Society* 104, 53-73. https://doi.org/10.1111/j.0066-7373.2004.00081.x

Krifka, Manfred, Pelletier, Francis J. Pelletier, Carlson, Gregory N., ter Meulen, Alice, Chierchia, Gennaro & Link, Godehard. 1995. Introduction to Genericity. In Gregory N. Carlson & Francis Jeffry Pelletier (eds.). *The Generic Book*, 1-124. Chicago: University of Chicago Press.

Krueger, Joachim & Clement, Russell. 1994. The truly false consensus effect: An ineradicable and egocentric bias in social perception. *Journal of Personality and Social Psychology*, 65, 596–610. https://doi.org/10.1037/0022-3514.67.4.596

Lasersohn, Peter. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy,* 28(6), 643–86. https://doi.org/10.1007/s10988-005-0596-x

Leslie, Sarah Jane, Khemlani, Sangeet & Glucksberg, Sam. 2011. Do ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language, 65,* 15-31. https://doi.org/10.1016/j.jml.2010.12.005

Moltmann, Friederike. 2010a. Generalizing Detached Self-Reference and the Semantics of Generic One. *Mind and Language* 25(4), 440-473. https://doi.org/10.1111/j.1468-0017.2010.01397.x

Moltmann, Friederike. 2010b. Relative Truth and the First Person. *Philosophical Studies* 150 (2), 187-220. http://www.jstor.org/stable/40856551.

Ninan, Dilip. 2014. Taste predicates and the acquaintance inference. *Semantics and Linguistic Theory (SALT*) 24, 290-309. https://doi.org/10.3765/salt.v24i0.2413

Nouwen, Rick. 2007. *Predicates of (im)Personal Taste.* Ms. University of Utrecht

Patel-Grosz, Pritty. 2012. *(anti-)Locality at the Interfaces*. Cambridge, Massachusetts: MIT PhD dissertation.

Pearson, Hazel. 2022. Individual and stage-level predicates of personal taste: another argument for genericity as the source of faultless disagreement. In Jeremu Wyatt, Julia Zakkou & Dan. Zeman (eds.), *Perspectives on Taste: Aesthetics, Language, Metaphysics and Experimental Philosophy,* 286-318. Routledge.

Ross, Lee, Greene, David, & House, Pamela. 1977. The "false consensus effect": an egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology* 13, 279-301. https://doi.org/10.1016/0022-1031(77)90049-X

Snyder, Eric. 2013. Binding, Genericity, and Predicates of Personal Taste. *Inquiry* 56 (2-3), 278-306. https://doi.org/10.1080/0020174X.2013.784485

Snyder, Melvin, Stephan Walter, & Rosenfield, David. 1978. Attributional egotism. In John H. Harvey, William Ickes & Robert F. Kidd (eds.), *New directions in attribution research*. Vol 2,91-117. Hillsdale, NJ: Lawrence Erlbaum.

Stephenson, Tamina. 2007. *Towards a theory of subjective meaning*. PhD, MIT.

Tessler, Michael H. & Goodman, Noah D. 2019. The Language of Generalization. *Psychological Review* 126(3), 395-436. https://doi.org/10.1037/rev0000142

Van Boven, Leaf, & Loewenstein, George. 2003. Social projection of transient drive states. *Personality and Social Psychology Bulletin* 29, 1159-1168. https://doi.org/10.1177/0146167203254597

Elsi Kaiser
Department of Linguistics
3601 Watt Way, GFS 301
University of Southern California
Los Angeles, CA 90089-1693
emkaiser@usc.edu

Haley Hsu
Department of Linguistics
3601 Watt Way, GFS 301
University of Southern California
Los Angeles, CA 90089-1693
haleyhsu@usc.edu