Proceedings of SALT 32: 000–000, 2022

Social identity, precision and charity: when less precise speakers are held to stricter standards.*

Andrea Beltrama
University of Pennsylvania

Florian Schwarz
University of Pennsylvania

Abstract Recent work has begun to show systematic connections between social information and pragmatic reasoning. Among other contributions, this work recently highlighted the role of constructs such as the *persona* embodied by the speaker in shaping the process with which comprehenders compute the extension of a numerical expression (Beltrama & Schwarz 2021). These findings, in turn, raise the question of whether similar effects obtain for other interpretation processes – in particular, when assessing the correctness of such imprecise descriptions in light of a single known and determined fact. We explore this question by testing the impact of speaker identity on T(ruth)-V(alue) J(udgment)s based on the interpretation of number words. We find that imprecise statements from speakers socially expected to be *less* precise – i.e. "Chill" ones – are rejected at a higher rate, and thus held to *more* stringent evaluation standards, than those from speakers socially expected to speak more precisely – i.e. "Nerdy" ones. This contrasts with the picture selection task results in Beltrama & Schwarz (2021), where utterances by Nerdy speakers are associated with a narrower standard of precision. We explain the new finding by appealing to the idea that, by virtue of generally being perceived to be more precise, Nerdy speakers are granted higher epistemic credibility than Chill ones. The emerging picture is one in which TVJ assessments are affected by social considerations in a different way from other experimental taskssuggesting a nuanced interplay between social information and different interpretation tasks and processes.

Keywords: Imprecision; social identity; numerals; truth-value judgments; personae; charity

1 Introduction

Much work in sociolinguistics and psycholinguistics has highlighted the central role of social information in language processing. In particular, research across

* We are grateful to a number of people for their feedback, comments, and guidance at different stages of this project. These include: Helena Aparicio, Alex Kalomoiros, Chris Kennedy, Lisa Miracchi, Gareth Roberts, Rista Plate, Hannah Read, Meredith Tamminga, Muffy Siegel, Alexis Wellwood, as well as fellow members of the Language and Cognition Lab and the Experimental Study of Meaning Lab at the University of Pennsylvania. All errors are our own.

different domains of speech perception and production – e.g., sound categorization and recognition; convergence; imitation – illuminated how comprehenders promptly utilize information about different properties of speaker. This information spans different layers of identity categories, including speakers' location of origin, gender, or race (Strand 1999; Niedzielski 1999; Hay, Warren & Drager 2006; Hay 2009; Babel 2012; Drager 2015; Staum Casasanto 2008; Sumner, Kim, King & McGowan 2014; Wade 2022), as well as the more idiosyncratic, specific "types" of person, or *personae*, that they embody.[1] For example, comprehenders primed with social types such as "Valley Girl" displayed different perceptions of vowel boundaries than listeners primed with other types (D'Onofrio 2015, 2018); similar effects of persona have been unveiled in connection with other aspects of speech processing, e.g., judgments of foreign accentedness (D'Onofrio 2019) or syntactic parsing (Choe, Sloggett, Yoshida & D'Onofrio 2019).

In a similar vein, a growing amount of work has been shedding light on the impact of social information on the processes involved when interlocutors resolve, and ultimately exchange, meaning in conversation. On an empirical level, it's been shown that pragmatic reasoning is impacted by a range of different social factors – e.g., politeness considerations (Bonnefon, Feeney & Villejoubert 2009; Yoon, Tessler, Goodman & Frank 2016, 2020; Mazzarella, Trouche, Mercier & Noveck 2018); affect (Kao, Wu, Bergen & Goodman 2014; Bergen 2016); speaker-specific properties such as linguistic nativeness (Fairchild & Papafragou 2018) and broader personae and identity categories (Beltrama & Schwarz 2021; see §2 for further detail). Conversely, it has also been shown that comprehenders promptly infer identity and personality features of a speaker from the semantic and pragmatic properties of their utterances (Acton & Potts 2014; Beltrama & Staum Casasanto 2017, 2022; Acton 2019; Glass 2015; Karawani & Waldon 2017; Jeong 2021; Thomas 2021; Hunt & Acton 2022; see Beltrama 2020 for an overview). In parallel with these empirical advances, general frameworks aiming to capture pragmatic inferences have been expanded in a variety of directions, reflecting a shift towards a more comprehensive approach to the study of the interpretation of linguistics utterances: for example, work within the *Rational Speech Act* framework has developed a system to model the interpretation of nonliteral meaning as an instance of social cognition more broadly – i.e., of rational action in which interlocutors aim at maximizing their utility in light of the goals that they are pursuing, similar to what happens in other domains of human behavior (Frank & Goodman 2012; Goodman & Stuhlmüller 2013; Kao et al. 2014; Lassiter & Goodman 2017; Goodman & Frank 2016). In a similar vein, recent developments of Gricean and game-theoretic models of communication productively

1 The category of *persona* has received extensive attention in sociolinguistics and linguistic anthropology, independent of its implications on language processing. See Irvine (2001); Agha (2005); Coupland (2007); Eckert (2008); Podesva (2011); Kiesling (2016) for foundational work in the area.

relied on the formal tools utilized in modeling meaning-related inferences to capture the processes of interlocutors signaling and negotiating social identity (Burnett 2017, 2019; Henderson & McCready 2019).

Taken together, these endeavors have led to a considerable expansion of the scope of semantics and pragmatics, highlighting novel questions around the dynamics whereby comprehenders recruit contextual information when interpreting meaning. Especially relevant, in this respect, is the issue of how considerations about speakers identity and personality shape comprehenders' pragmatic reasoning with respect to phenomena that had long been investigated before, but that had rarely been considered under a view that takes into account the distinctive social profiles of the interlocutors. The present paper takes a step towards illuminating this issue by exploring the following question: how do the distinctive identity features of a speaker impacts comprehenders' determination of whether, given a known fact, a certain linguistic description appropriately represents this fact? Besides representing a task that interlocutors need to routinely navigate in everyday communicative situations, this type of assessment also plays a central methodological role in experimental semantics and pragmatics, as it stands at the core of the standard Truth Value Judgement (TVJ) task (henceforth, **TVJ**; Crain & McKee 1985; Crain 1998) – a paradigm which has been extensively used for studying comprehenders' interpretation processes (Noveck 2001; Papafragou & Musolino 2003; Bott & Noveck 2004 i.a.). In particular, even though recent work has begun to explore how respondents reason about the discourse situation and their relation to it to make these assessments (Sikos, Kim & Grodner 2019; Waldon & Degen 2020; Scontras & Pearl 2020), the role of social information in determining the outcome of these assessments remains largely uncharted, opening up an intriguing question that could have important theoretical and methodological ramifications for the study of meaning.

In this paper, we shed light on this issue by focusing on a case study that emerges as an ideal testbed for our purposes: imprecision resolution in the interpretation of numerical expressions – a pragmatic phenomenon which already has been investigated in other interpretation task paradigms with regards to the impact of persona-based information on comprehension processes (Beltrama & Schwarz 2021; see §2). Specifically, we set out to test how the persona embodied by a speaker – that of a *Nerdy* vs. a *Chill* person – shapes the way in which comprehenders judge the correctness of the imprecise statements uttered by this speaker – e.g., "The price is $200" when the price is in fact $207. As discussed in §3, our findings indicate that imprecise descriptions uttered by Chill speakers, who are expected to characterize facts less precisely, are rejected at a higher rate than imprecise descriptions uttered by Nerdy speakers, who are instead expected to characterize facts more precisely. We take this result to support two conclusions. First, when making assessments in a TVJ paradigm, comprehenders do indeed track information about the speaker

social profile; second, the way in which they reason about this information in this task paradigm seems to differ in interesting ways from the way it is considered in other paradigms, such as the one employed by Beltrama & Schwarz (2021), where comprehenders effectively draw inferences in the opposite direction to the one tested in a TVJ setup – i.e., they are asked to infer what facts could be taken to go along with an utterance (see Beltrama & Schwarz 2021 and §2 for further details). All in all, we believe these results highlight a novel aspect in which interpretive judgments are influenced by social information, and thus contribute to shedding further light on the context sensitivity of a task that remains very much central in semantics and pragmatics and, more broadly, on the intricate interaction between social information and pragmatic reasoning in interpretation.

The paper proceeds as follows. §2 introduces (im)precision and the relevant prior research on the phenomenon; §3 reports on an experiment that we carried out to investigate our question; §4 presents a general discussion; §5 concludes.

2 Imprecision: a case study

Numerical expressions, such as those in (1), are known for having context-sensitive meanings.

(1) a. It's 6 o'clock.
b. The ticket costs $300.

While times and prices can be intuitively be seen as denoting precise values, they are often used in a looser way: for instance, one could use (1-a) to describe a time of 5:57; or (1-b) to describe a price of $295. This phenomenon, known as *imprecision*, has been extensively investigated across philosophy and linguistics (Austin 1962; Lewis 1979; Pinkal 1995; Lasersohn 1999; Syrett, Kennedy & Lidz 2009; Kennedy 2007; Morzycki 2011; Burnett 2014; Solt 2014; Klecha 2014; Cummins, Sauerland & Solt 2012; Kao et al. 2014; Aparicio 2017; Beltrama & Hanink 2019; Beltrama 2021): central to this work is the idea that it is indeed possible for speakers to use numerical expressions in a way that isn't fully adherent to the facts being described – i.e., by producing statements that are literally false, but "close enough" to being true (Lasersohn 1999).[2]

Two distinctive properties of imprecision make this phenomenon suitable for our purposes. One is that the possibility of speaking imprecisely introduces a space of *indeterminacy* in the interpretation of numerals, which requires listeners to engage in inferential work (Van Der Henst, Carles & Sperber. 2002; Cummins et al. 2012; Solt, Cummins & Palmovic 2017; Aparicio 2017; Syrett & Aravind 2021; Krifka 2007;

2 See Solt (2014) and Sauerland & Stateva (2011) for a different take.

Kennedy 2007; Leffel, Xiang & Kennedy 2016): because the appropriate amount of imprecision that can be applied in a communicative scenario is not encoded in the lexical meaning of a numeral, comprehenders need to reason about a variety of contextual cues to ultimately determine what this threshold is, and adjust their interpretive behavior on a case by case basis (see Van Der Henst et al. 2002; Kao et al. 2014; Solt et al. 2017; Aparicio 2017 for work on the cognitive underpinnings of this reasoning).

The other distinctive property of (im)precision is that, as it has already been shown in previous work, comprehenders' reasoning about imprecision is closely connected to their reasoning about the speaker's distintive identity features. On the one hand, listeners reason rely on precision as a cue to draw social inferences about speaker. For example, speakers using sharp numbers (e.g., "207") – normally taken to signal a high level of precision (Krifka 2007) – are associated with a cluster of social qualities that can be traced back to two core dimensions of social evaluation: high intellectual status – e.g., as articulate, educated, intelligent; and low sociability – e.g., as annoying, pedantic, obsessive, and generally unlikable (Beltrama 2018; Beltrama, Solt & Burnett 2022).[3] Similarly, speakers using numerals in an explicitly imprecise fashion (i.e., "around 200") are perceived as embodying opposite sets of qualities – e.g., as friendlier and more laid-back, but less intelligent and educated (Beltrama et al. 2022). On the other hand, and even more centrally for our purposes, it's been shown that the social information available in the context affects the computation of the level of precision with which a numeral is interpreted. In particular, Beltrama & Schwarz (2021) provide evidence that comprehenders reason about the persona embodied by the speaker in inferring the range of facts compatible with a numeral description. Evidence supporting this claim is based on participants' *indirect* judgments about whether a certain fact is a likely candidate for being the one the numeral utterance was about. Specifically, in a picture selection paradigm called the COVERED SCREEN TASK[4], participants were presented with a conversation between one of two sets of characters, manipulated between subjects: the Nerdy characters, who embody the social features generally associated with precise speech; or the Chill characters, who instead embody the social features generally associated with less precise speech. The two sets of characters, which had been previously normed to ensure that they would be perceived as socially different, and indeed associated with different expectations in terms of how precisely they speak, are shown in Figure 1.[5]

3 See Welsh, Navarro & Begg (2011); Mason, Lee, Wiley & Ames (2013); Xie & Kronrod (2012); Zhang & Schwarz (2011) for similar findings from the social psychology and marketing literature.

4 The task was a variant of the "Covered Box Paradigm" used in experimental studies on meaning. See in particular Huang, Spelke & Snedeker (2013).

5 See Beltrama & Schwarz (2021) for further details on the norming study.

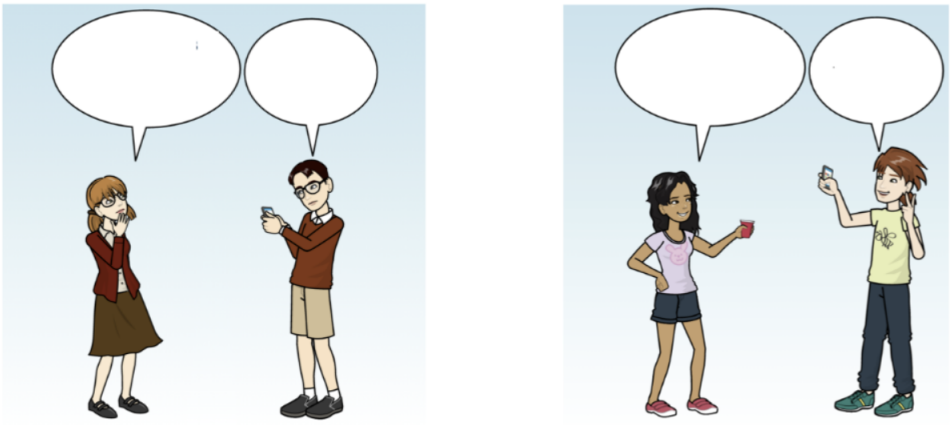

Figure 1 Nerdy (left) vs. Chill (right) characters

In the experimental task, participants would see one of these two sets of characters engaging in a conversation, which always concluded with one character uttering a numeral description (e.g., “It’s $200"). In the critical condition, they were then shown a picture of a screen showing a slightly mismatching number (“$207"; **VISIBLE** choice), and were instructed to select this picture if they thought it fit the content of the utterance, indicating an imprecise interpretation; as an alternative, they could select a screen turned face down (**COVERED** choice), indicating a rejection of the mismatching number, and thus a more precise interpretation. The experimental set-up is shown below, with the Nerdy condition for the Persona manipulation.

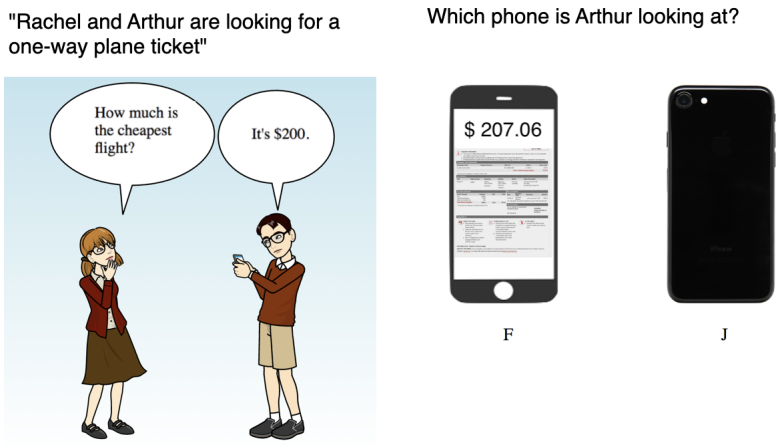

Figure 2 Display before making the choice (Persona Condition: Nerdy)

As predicted, respondents turned out to be more inclined to accept a slightly mismatching number as being correctly described by a numeral when the speaker embodied a Chill, as opposed to a Nerdy persona; this suggests that comprehenders reason about the social identity of the speaker when determining the margin of imprecision that can be associated with a numeral in a given context – and thus adjust the pragmatic extension of the numeral accordingly.

But these findings also raise the question of how the interplay between social information and the relevant pragmatic reasoning plays out in reverse interpretation tasks: that is, does the perception of speaker persona similarly affect numeral interpretation when comprehenders are *directly* asked to determine whether an utterance correctly describes a given fact? Put it differently: given an imprecise statement, can the speaker's description can still be accepted as being close-enough to the facts, or should instead be deemed a mischaracterization? This type of interpretative judgment is central to both many communicative situations and common experimental tasks; as a result, we believe it is important to investigate how social identity factors shape comprehenders' reasoning about this specific pragmatic dimension as well. Besides allowing for a more comprehensive understanding of how (im)precision is resolved, exploring this question also affords, on a broader level, the possibility of exploring what is thus far an uncharted domain of the interplay between social information and linguistic interpretation, providing a novel angle to consider how social information shapes a type of inferences that, as discussed in §1, stands at the core of the Truth Value Judgment paradigm: an experimental approach that has been extensively utilized in the study of meaning, but whose sensitivity to social considerations remains largely unexplored.

As for the nature of this effect, we predict that speaker persona could affect the comprehenders'assessment of whether an imprecise statement can be taken to correctly describe a particular fact in one of two possible ways. One possibility is that comprehenders assess the acceptability of an imprecise numeral by reasoning about the range of values that the numeral could plausibly be mapped onto in the context, parallel to the Covered Screen task in Beltrama & Schwarz (2021). If this is the case, imprecise descriptions produced by Nerdy speakers, which previous work showed to be associated with a more narrow range of values, should be rejected as mischaracterizations *more* often than imprecise descriptions produced by Chill speakers. We refer to this possibility as **Hypothesis 1**. But an alternative possibility also exists: comprehenders might determine whether they intend to accept or reject an imprecise statement by recruiting social information to reason on the epistemic level, rather than the descriptive one – that is, by leveraging social information to assess the amount of *credibility* that they are willing to ascribe to a speaker producing an imprecise statement. Accordingly, comprehenders could posit that, because Nerdy speakers are generally more accurate in their descriptions, they likely have some valid

reason to be imprecise in the specific context; as a result, they are more deserving to being given the benefit of the doubt in the case of an imprecise statement than speakers embodying the qualities linked to lower precision. This should lead to the opposite pattern to the one outlined above: imprecise descriptions produced by Nerdy speakers should be rejected *less* often – i.e., accepted as appropriate to being true more often – than imprecise descriptions produced by approximate speakers. We refer to this possibility as **Hypothesis 2**.

To test these hypotheses, we implement a TVJ version of the COVERED SCREEN paradigm, creating a task in which respondents are presented with both a description and an established body of facts, and are asked to assess the appropriateness of the former to represent the latter. We now proceed to discuss this study in §3.

3 The experiment

We implemented this setup by leaving the design of Beltrama & Schwarz (2021)'s COVERED SCREEN task described above unchanged except for one crucial modification: following the conversation between the Nerdy or Chill characters, participants were only shown the picture of one phone, with a number visible on the screen. They were told that that phone was the one that the speaker was indeed looking at, and asked to indicate whether the character's response was RIGHT or WRONG, with the former response taken to represent an acceptance of the utterance, and the latter a rejection thereof.

3.1 Methods & Design

Our stimuli utilized visually displayed dialogues, like those Beltrama & Schwarz (2021), with variants resulting from manipulations that crossed two factors in a 2×3 design.

Our first manipulation, implemented between-subjects, varied the persona embodied by the displayed characters (Nerdy: Arthur and Rachel vs. Chill: Alex and Eva). Each set of characters was framed in a conversation prefaced by a brief context sentence; in the dialogue, the female character and the male character would respond based on information they accessed by looking at their phone, uttering a quantity expression in the form of a round number. The presentation was identical to the cartoons in Figure 1, with the question and the response provided as text in the speech bubble.

The second manipulation was a within-subjects factor. After seeing the dialogue between the characters, participants were asked the question "Which phone is Arthur/Alex looking at" and were shown the picture of one phone, with a number visible on the screen. They were told that that phone was the one that the character

was indeed looking at, and asked to indicate whether the character's response was RIGHT or WRONG in light of the displayed number. The number displayed on the phone was manipulated depending on its relation to the one uttered by the character, with three possible levels: *Match* with identical displayed and uttered numbers; *Mismatch*, with a large divergence between the two; and the critical *Imprecise* level, with only a slight divergence between the uttered and the displayed number. An example from the three Screen Fit levels for a given item is provided in Figure 3.

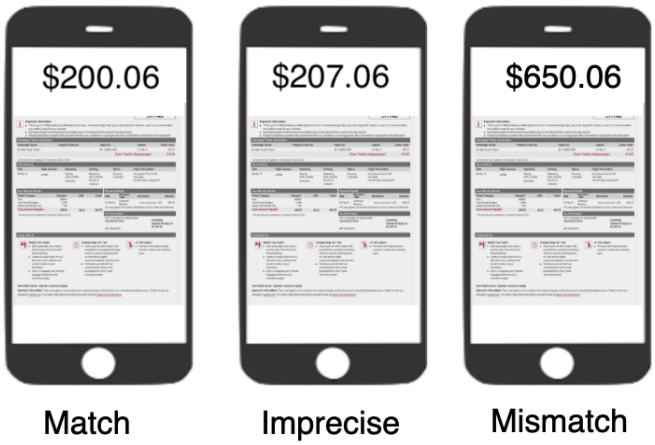

Figure 3 Screen Fit Manipulation

The full display of a item is shown in Figure 5, with Nerdy as the level for the Persona manipulation, and Imprecise for the Screen Fit condition.

3.2 Materials

24 experimental items were created, each varied across 6 different conditions resulting from the 2×3 manipulation of the factors described above. The Persona manipulation was administered between-subjects: a given participant was either assigned to dialogues between the Nerdy characters or between the Chill characters. The Screen Fit manipulation was administered within-subjects: each participant saw 6 items in the Match and the Mismatch conditions and 12 items in the Imprecise condition, with item-condition pairings counterbalanced in a Latin Square Design. 8 items contained utterances describing prices, expressed in dollars (as in Figure

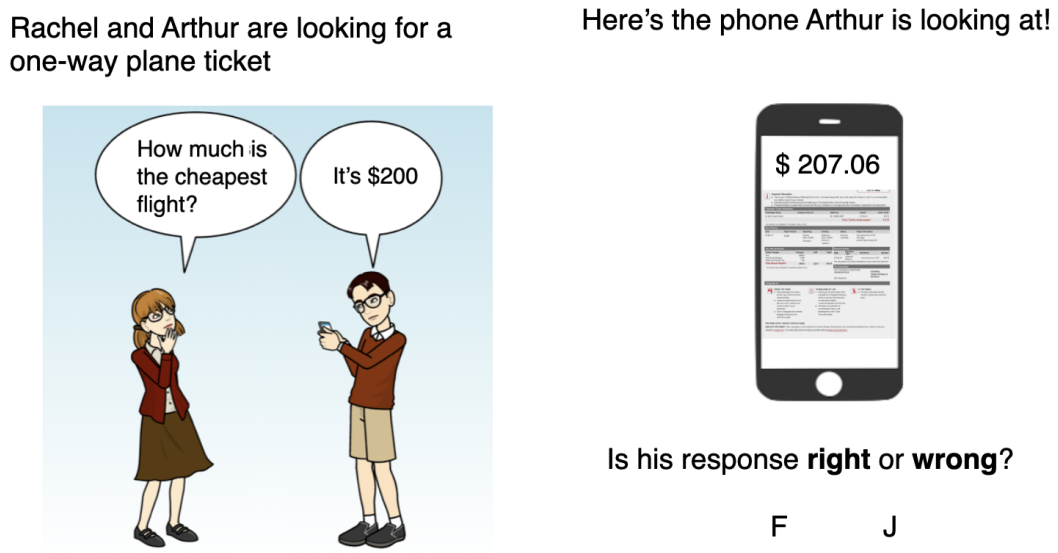

Figure 4 Display before making the choice (Condition: Nerdy, Imprecise)

3-4); 8 items contained utterances describing distances, expressed in miles; and 8 items contained items describing times, expressed in hours and minutes. Across the three types of numerals, the range of deviation in the Imprecise condition was always comprised between 5% and 18%: cost and distance diverged from the uttered value by 5, 6, 8, 9, 12, 14, 16, or 18; and time diverged by 1, 2, 3, 4, 5, 7, 9, 11, or 13. The ranges of divergence for times were smaller since proportionally, a minute makes up a greater proportion of an hour (1/60th) than 1 mile or dollar makes of a 100 miles/dollars, and quarter hour intervals intuitively constitute a salient degree of coarseness.

The experiment also included 24 filler items, which featured a dialogue between two separate characters called Becky and Tyler. The dialogue was also prefaced by a brief context sentence and was concluded by Tyler uttering a description with the quantifier *some*. In 8 filler items the visible screen would report a list of items which would make the description true; in 8 filler items it would report a list of items that would make the description false; and in 8 filler items, it would make the description under-informative, in that an 'all' statement would also have been true, thus yielding a condition where acceptance of the visible picture would correspond to adopting a literal 'some and possibly all' meaning of 'some', parallel to standard studies on scalar implicatures. An example of a filler item is shown below.

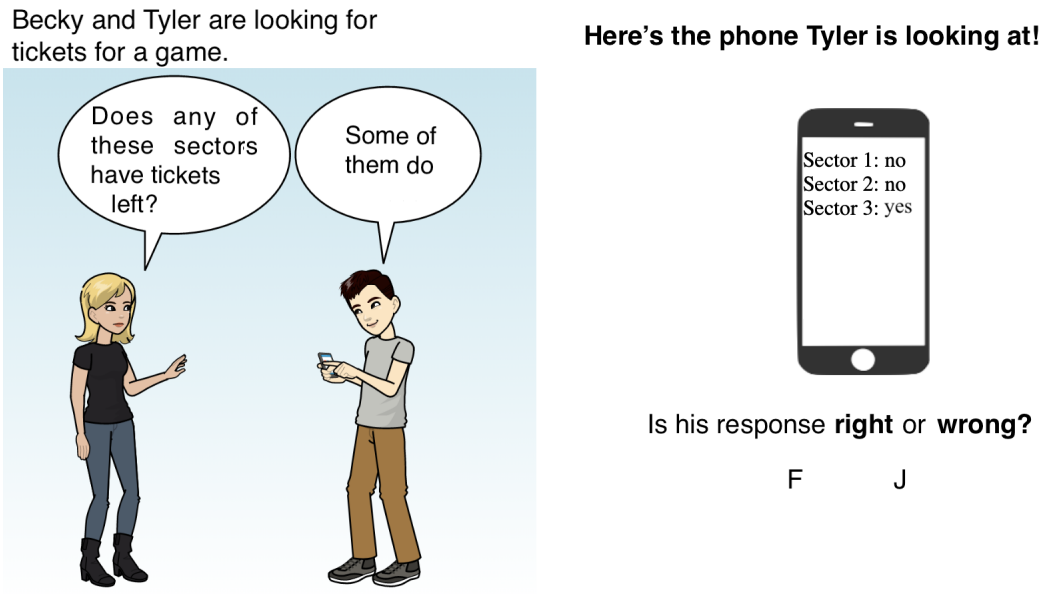

Figure 5 Full filler item

The fillers were alternated with the experimental items, so that participants would never see two consecutive occurrences of a filler or an experimental item.

3.3 Procedure

The study was implemented and administered online on the PCIbex platform[6] (see Schwarz & Zehr 2021 for details). For each item, the context sentence was introduced first on the top left of the screen. After a 4-second pause to provide time to read this, images of the two characters would appear; after another 3-second pause to allow time to look at these characters (and the critical visual cues on their persona), the question asked by the first character was shown in the form of a speech bubble; following another 3 second pause, the answer from the other character was displayed to complete the dialogue. Finally, the question reminding the participant of the experimental task as well as the two picture of the phone appeared on the right-hand side of the screen. Participants entered their responses by pressing the key matching the letter displayed under the picture on the keyboard. The experimental items were preceded by three practice filler items in which the response would involve the use of a quantifier, and the content of the visible screen would be either a perfect match

6 https://www.pcibex.net

or an obvious mismatch. Feedback was provided on the practice items, so as to help the participants familiarize themselves with the task.[7]

3.4 Participants

190 participants were recruited on Prolific and compensated $2 for participating (Age range: 18-66; Age Mean: 24; female = 132; male = 53; other = 5). The median duration of the study was 15 minutes ($8/hour).

3.5 Results

As can be seen in Figure 6, rejections (=“wrong" responses) were at ceiling and at floor in the Mismatch and Match condition respectively, with intermediate, mid-range rejection rates in the Imprecise condition, as in Experiment 1. The differences between the Imprecise condition and the controls was again highly significant in mixed effect logistic regression models (Match: β=-10.06; SE=0.73; $z = 13.72$, $p < 0.001$); Mismatch: β=6.68; SE=0.44; $z = 15.03$, $p < 0.001$). To assess the effect of Persona in the critical, Imprecise condition, we fit a mixed-effects logistic regression with Persona as fixed effect, and by-Subject and by-Item random intercepts.[8] The rejection rates for the two Persona levels differed significantly from each other, with higher rejection rates for Chill speakers than for Nerdy speakers (β=1.13; SE=0.52; $z = 2.18$, $p < 0.05$).

7 Following the last trial, all participants, regardless of whether they had been assigned to the Nerdy or the Chill condition, were asked to complete the two question exit questionnaire shown below:

(i) a. I’d describe myself as: 1=not chill at all; .... 10=very chill
b. I’d describe myself as: 1=not nerdy at all; ... 10=very nerdy

The aim of the questionnaire was to assessing the degree to which they saw themselves as sharing key qualities of the speakers in the experiment. For reasons of space we will not be discussing these data.

8 Again, this model constituted the maximally complex random effect structure that would converge.

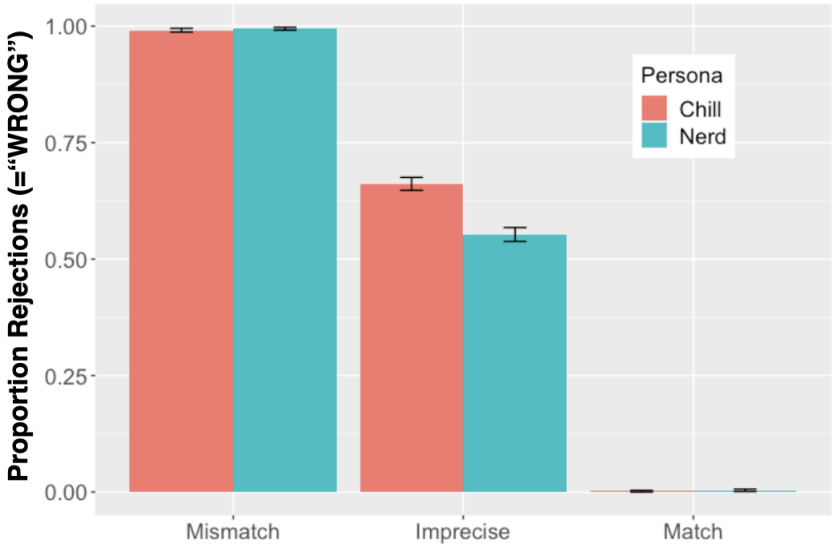

Figure 6 COVERED choices across Screen Fit, split by Persona

4 Discussion

The findings from our study indicate that speakers' social identity affects comprehenders' assessments of imprecise statements: the same numerical descriptions are rejected – i.e. judged as WRONG – at a higher rate when produced by Chill speakers than by Nerdy ones. Combined with those from prior work, these results provide further evidence that comprehenders recruit social information about the speaker when engaging in pragmatic reasoning – and do so across different tasks and situations. At the same time, the specific impact of a speaker's identity and personality features on the interpretation of what they say appears to crucially depend on the specific type of interpretive judgment and corresponding pragmatic reasoning involved in a given task context. Specifically, when comprehenders have to decide whether a candidate fact is part of the range described by a numeral in context, they take numerals uttered by Nerdy speakers to have a more narrow extension than numerals uttered by Chill speakers – effectively applying *stricter* standards of precision for the former than the latter (see results from Beltrama & Schwarz 2021). But when comprehenders have to assess an imprecise numeral statement relative to a fixed, given fact, as in the current study, they are more likely to accept imprecise statements by Nerdy speakers than by Chill ones – effectively being more *lenient* towards the former than the latter. This suggests that when it comes to assessing the suitability of a description, comprehenders recruit social information to reason on the epistemic level, as opposed to the descriptive one – that is, to calibrate the amount of credibility

that they are willing to ascribe to the speaker. More specifically, we suggest that comprehenders reasoned about the fact that, even though the numerals uttered by Nerdy speakers are associated with a stricter range of possible values as referents, Nerdy speakers are generally perceived to be more accurate in their descriptions than Chill ones. As a result, they are treated by comprehenders as more credible sources in situations presenting a discrepancy between a description and a state of affairs, leading to a higher rate of RIGHT responses, in line with **Hypothesis 2** outlined in §2.

Taken together, these results suggest that considerations about the persona embodied by the speaker shape how comprehenders' form an assessment of the appropriateness of a given description to represent a particular fact – and thus affect the outcome of the response behavior in a Truth Value Judgment task. Before considering the broader implications of these findings, however, it is important to first address an obvious question that is raised by the results from the experiment: *why* was this particular pattern observed, as opposed to the alternative one outlined in Hypothesis 1 (see §2) – i.e., that Nerdy speakers' descriptions, by virtue of being associated with a more narrow extension, should actually be rejected more often? We propose that this result can be illuminated by considering the distinct epistemic implications linked to rejecting a description across the two paradigms. In Beltrama & Schwarz's Covered Screen Task, rejecting the imprecise display and selecting the COVERED option does not amount to an evaluative judgment of the speaker: it simply indicates that comprehenders believed that the utterance had been made in the presence of a more closely fitting state of affairs than the one shown in the VISIBLE display – a determination still compatible with taking the speaker to be cooperative, truthful, and overall conversationally competent. By contrast, a rejection in a TVJ task is crucially *prejudicial*: it commits the comprehender to holding the speaker *blameworthy* for not abiding to the norms of conversation – at the very least, on account of describing things inaccurately, and thus using language untruthfully. We suggest that it is precisely the prejudicial nature of a rejection in the current experiment that tipped the scales against the response pattern predicted by Hypothesis 1, and in favor of the one we actually observe. Because Nerdy speakers are generally (perceived to be) more accurate in their descriptions, they are presumably taken to be less likely to violate the relevant conversational norms than Chill ones. Accordingly, higher WRONG response rates for Chill speakers can be explained by positing that their stereotypical representation as imprecise language users made them more prone than Nerds to being seen as guilty of violating relevant conversational norms. This reconciles the seemingly conflicting results from the two studies: the expectation of higher accuracy for Nerdy speakers indeed goes hand in hand with a higher level of credibility assigned to them, leading to a more charitable disposition towards these speakers when it comes to assessing their statements in light of a given fact.

All in all, these results enrich our understanding of how social information shapes processes of meaning interpretation and meaning ascription. As summarized in §1, a growing number of studies have illuminated this issue (see Beltrama 2020 for an overview). This work provided important contributions both at the level of empirically demonstrating the interaction between social considerations and pragmatic reasoning (Bonnefon & Villejoubert (2006); Bonnefon et al. (2009); Acton & Potts (2014); Beltrama & Staum Casasanto (2022); Mazzarella et al. (2018); Fairchild & Papafragou (2018); Mahler (2020), and of incorporating social information in broader communicative frameworks (Goodman & Frank 2016; Yoon et al. 2016; Kao et al. 2014; Bergen 2016; Burnett 2017, 2019; Henderson & McCready 2019; Acton 2019). Our findings broaden the scope of such endeavors by highlighting personae as a source of contextual information shaping how comprehenders compute context-sensitive parameters – e.g., the (im)precision threshold – that are central to determining whether a particular description adequately represents a state of affairs – a crucial evaluation that interlocutors must make whenever interpreting a meaning of a sentence. We take this conclusion to carry two important ramifications for the study of meaning.

One is that the reasoning behind the interpretation of numerals turns out to share a common denominator with other, seemingly distinct domains of reasoning – both within and beyond the linguistic domain. Looking at language processing, it has been long known that different realms of speech perception are heavily influenced by the available social information about the speaker, both at the level of a speaker's demographic profile (Strand 1999; Niedzielski 1999; Hay et al. 2006; Hay 2009; Babel 2012; Drager 2015; Staum Casasanto 2008; Sumner et al. 2014; Wade 2022), and at the level of the more holistic constructs salient in the context, such such as the stereotypical *persona* embodied by conversational participants. While sociolinguistic theory has long highlighted these constructs as central to how identity is signaled and perceived in interaction (Eckert 1989; Irvine & Gal 2000; Bucholtz 2001; Zhang 2005; Agha 2005; Podesva 2011; Kiesling 2016; King 2021), it is not until recently that persona-level information has also been shown to play a central role to processes of language comprehension and perception (D'Onofrio 2015, 2019; Choe et al. 2019; see D'Onofrio 2020 for an overview). Combined with our results from prior work, the findings from this study crucially suggest that parallel effects of persona-based information are also observed in processes of linguistic interpretation: in particular, they pinpoint comprehenders' social perception of speakers as shaping the inference-drawing mechanisms whereby the literal content of an utterance is enriched with contextual information; and they highlight this process as one that, similar to other domains of language processing, crucially hinges on how comprehenders perceive the identity and personality features of the conversational agents involved in the exchange.

The other important implication of our findings is a methodological one. In particular, we have provided novel evidence that comprehenders' social perception of the speaker shapes the outcome of judgments of descriptive appropriateness in a TVJ task – a behavioral measure widely utilized in the experimental study of meaning, and commonly seen as a window into the outcome of semantic interpretation (Noveck 2001; Papafragou & Musolino 2003; Bott & Noveck 2004; see §1). In recent years, a growing body of work called for a critical re-assessment of the nature and significance of these judgments, highlighting how they are often affected by factors – and thus sources of variability – that go beyond the sheer linguistic properties of a given utterance. These include, among others, the discourse context (Sikos et al. 2019; Scontras & Pearl 2020); the probability with which respondents see themselves as producing the utterance (Waldon & Degen 2020); or properties inherent to the demographic background of the speaker, such as their linguistic nativeness (Fairchild & Papafragou 2018; Fairchild, Mathis & Papafragou 2020). Our results open a novel perspective on the context-sensitivity of this behavioral measure, suggesting two takeaways. First, when making these assessments, comprehenders track an array of contextual factors that extend beyond those traditionally investigated in pragmatics, and include highly specific social constructs such as the speaker persona. Second, the way in which these factors impact comprehenders' determination of the correctness of an utterance in a TVJ task does not align with the way in which the same factors affect other interpretation tasks on the same phenomenon – e.g., the COVERED SCREEN picture selection task utilized in Beltrama & Schwarz (2021). Especially noteworthy, here, is the observation that, in a TVJ setup, the perceived appropriateness of an imprecise description is *inversely* correlated with the expectation that a speaker uses this description imprecisely; and that, more broadly, the interplay between social reasoning and TVJ assessments cannot be fully understood without considering the interplay between a speaker's social identity and the prejudicial implications entailed by rejecting their utterance – and in particular, the ascription of conversational blameworthiness that goes hand in hand with deeming an utterance inappropriate. We see this interaction as one that remains crucially underexplored in semantic and pragmatic work, calling for a more extensive consideration of how TVJ assessments, and other experimental measures deployed in the study of meaning, are affected by information about both the identity of who produces a particular utterance, and the social implications entailed by choosing a particular response choice as opposed to another.

Looking at the broader picture, it is important to consider how the effects of persona on meaning interpretation observed in our study relates to how social information shapes reasoning and action outside the linguistic domain proper. In particular, work on *testimonial injustice* in philosophy has suggested that people recruit social information about their interlocutors to determine the degree of authority

and knowledge that they are willing to ascribe to them on a given matter; and that this reasoning often results in prejudicial practices such as a *credibility deficit* (Fricker 2007) – observed, for example, when someone is treated as lacking authority on a matter by virtue of being a woman; or *credibility excess* (Davis 2016) – observed for instance when someone asks an Asian-American seated nearby to help them with their math problem, relying on the flawed assumption that everyone embodying this social identity is proficient in the discipline. In this respect, our findings highlight a parallel dynamics with respect to assessing imprecise descriptions: when faced with the task of determining the correctness of an imprecise statement, comprehenders ultimately grant more credibility to Nerdy speakers than to Chill ones, suggesting that, similar to the cases discussed in the testimonial injustice literature, comprehenders recruit social information to make decisions that can be highly impactful on how authoritative, credible or trustworthy a particular conversational agent is seen as in the context. Framed in the broader picture of how social stereotypes shape people's behavior, our results thus suggest that interpretation processes are not just subject to the effect of social information, but are also likely to be subject to, and simultaneously reproduce, similar patterns of bias and prejudice to those observed in other domains of human action. We see this as opening up yet another uncharted dimension for the study of linguistic interpretation: similar mechanisms linking social identity, meaning interpretation and the ascription of credibility are indeed likely to be at play – and possibly even more impactful – when more socially prejudicial dimensions of identity than Nerdy or Chill are involved – e.g., race, gender socioeconomic status – in line with what has been suggested in work on linguistic discrimination more broadly (see Jones, Kalbfeld, Hancock & Clark 2019 among many others). This, in turn, highlights the importance of engaging in further work aimed at unpacking the link between processes of meaning comprehension and the perception and evaluation of social identity, so as to better understand the interplay of these two domains of reasoning in people's communicative behavior.

5 Conclusion

We presented experimental evidence showing that the social persona embodied by the speaker impacts the process whereby comprehenders assess the appropriateness of imprecise descriptions to represent a given fact. Together with the findings from the growing amount of work investigating the interaction between social information and pragmatic reasoning, these results advance our understanding of the contextual factors that underlie meaning interpretation; and thus contribute to enriching semantics and pragmatics with insights from the study of how speakers' signal, negotiate and reason about their interlocutors' social identity in communication – a domain of research that has only been tangentially relevant to the study of meaning thus far.

References

Acton, Eric. 2019. Pragmatics and the social life of the english definite article. *Language* 95(1). 37–65. doi:10.1353/lan.2019.0010.

Acton, Eric & Chris Potts. 2014. That straight talk. Sarah Palin and the sociolinguistcs of demonstratives. *Journal of Sociolinguistics* 18(1). 3–31. doi:https://doi.org/10.1111/josl.12062.

Agha, A. 2005. Voice, footing, enregisterment. *Journal of Linguistic Anthropology* 15. 38–59.

Aparicio, Helena. 2017. *Processing context-sensitive expressions: the case of Gradable Adjectives and Numerals*: University of Chicago PhD dissertation.

Austin, John L. 1962. *How to do things with words*. Cambridge: Harvard University Press.

Babel, Molly. 2012. Evidence for phonetic and social selectivity in spontaneous phonetic imitation. *Journal of Phonetics* 40(1). 177–189. doi:10.1016/j.wocn.2011.09.001.

Beltrama, Andrea. 2018. Precision and speaker qualities. the social meaning of pragmatic detail. *Linguistics Vanguard* 4(1). doi:https://doi.org/10.1515/lingvan-2018-0003.

Beltrama, Andrea. 2020. Social meaning in semantics and pragmatics. *Language and Linguistics Compass* 14(9). doi:10.1111/lnc3.12398.

Beltrama, Andrea. 2021. Just perfect, simply the best: an analysis of emphatic exclusion. *Linguistics and Philosophy* doi:10.1007/s10988-021-09326-x.

Beltrama, Andrea & Emily Hanink. 2019. Marking imprecision, conveying surprise: *Like* between hedging and mirativity 55(1). 1–34. doi:10.1017/S0022226718000385.

Beltrama, Andrea & Florian Schwarz. 2021. From personae to interpretation: How social meaning affects pragmatic reasoning. Proceedings of Semantics and Linguistics Theory (SALT) 31 (forthcoming).

Beltrama, Andrea, Stephanie Solt & Heather Burnett. 2022. Context, precision and social perception: a socio-pragmatic study. In press at *Language in Society*.

Beltrama, Andrea & Laura Staum Casasanto. 2017. Totally tall sounds totally younger. intensifiers at the socio-semantic interface. *Journal of Sociolinguistics* 21(2). 154–182. doi:https://doi.org/10.1111/josl.12230.

Beltrama, Andrea & Laura Staum Casasanto. 2022. The social meaning of semantic properties. In Lauren Hall-Lew, Emma Moore & Robert J. Podesva (eds.), *Social Meaning and Linguistic Variation: Theorizing the Third Wave.*, Cambridge University Press. doi:10.1017/9781108578684.004.

Bergen, Leon. 2016. *Joint inference in pragmatic reasoning*: Massachusetts Institute of Technology PhD dissertation.

Bonnefon, Jean-FranÃ§ois, Aidan Feeney & GaÃ«lle Villejoubert. 2009. When some is actually all: Scalar inferences in face-threatening contexts. *Cognition* 112(2). 249 – 258. doi:https://doi.org/10.1016/j.cognition.2009.05.005.

Bonnefon, Jean-François & Gaëlle Villejoubert. 2006. Tactful or doubtful?: Expectations of politeness explain the severity bias in the interpretation of probability phrases. *Psychological Science* 17(9). 747–751. doi:10.1111/j.1467-9280.2006.01776.x. PMID: 16984289.

Bott, Lewis & Ira A. Noveck. 2004. Some utterances are underinformative: The onset and time course of scalar inferences. *Journal of Memory and Language* 51(3). 437–457. doi:10.1016/j.jml.2004.05.006.

Bucholtz, Mary. 2001. The whiteness of nerds: Superstandard English and racial markedness. *Journal of Linguistic Anthropology* 11. 84–100. doi:https://escholarship.org/uc/item/0663j2rg.

Burnett, Heather. 2014. From quantification and intensification to slack regulation: the case of adjectival ALL. In *Proceedings from the 48th Meeting of the Chicago Linguistic Society*, .

Burnett, Heather. 2017. Sociolinguistic interaction and identity construction: The view from game-theoretic pragmatics. *Journal of Sociolinguistics* 21(2). 238–271. doi:10.1111/josl.12229. https://onlinelibrary.wiley.com/doi/abs/10.1111/josl.12229.

Burnett, Heather. 2019. Signalling games, sociolinguistic variation and the construction of style. *Linguistics and Philosophy* 42. 419–450. doi:https://doi.org/10.1007/s10988-018-9254-y.

Choe, June, Shayne Sloggett, Masaya Yoshida & Annette D'Onofrio. 2019. Personae in syntactic processing: Socially-specified agents bias expectations of verb transitivity. Poster presented at CUNY 2019.

Coupland, Nick. 2007. *Style: Language Variation and Identity*. Cambridge University Press.

Crain, Stephen. 1998. *Investigations in universal grammar : a guide to experiments on the acquisition of syntax and semantics*. Cambridge, Mass: MIT Press.

Crain, Stephen & Cecile McKee. 1985. Acquisition of structural restrictions on anaphora. *North East Linguistics Society* 16(1).

Cummins, Chris, Uli Sauerland & Stephanie Solt. 2012. Granularity and scalar implicature in numerical expressions. *Linguistics and Philosophy* 35. 135–169.

Davis, Emmalon. 2016. Typecasts, tokens, and spokespersons: A case for credibility excess as testimonial injustice. *Hypatia* 31(3). 485–501. doi:https://doi.org/10.1111/hypa.12251. https://onlinelibrary.wiley.com/doi/abs/10.1111/hypa.12251.

D'Onofrio, Annette. 2015. Persona-based information shapes linguistic perception: Valley girls and california vowels. *Journal of Sociolinguistics* 19(2). 241–256.

doi:https://doi.org/10.1111/josl.12115.

D'Onofrio, Annette. 2018. Personae and phonetic detail in sociolinguistic signs. *Language in Society* 47(4). 513–539. doi:10.1017/S0047404518000581.

D'Onofrio, Annette. 2019. Complicating categories: Personae mediate racialized expectations of non-native speech. *Journal of Sociolinguistics* 23(4). 346–366. doi:10.1111/josl.12368.

D'Onofrio, Annette. 2020. Personae in sociolinguistic variation. *WIREs Cognitive Science* 11(6). doi:10.1002/wcs.1543.

Drager, Katie K. 2015. *Linguistic variation, identity construction and cognition*. Language Science Press. doi:10.17169/LANGSCI.B75.22.

Eckert, Penelope. 1989. *Jocks and Burnouts: Social Identity in the High School*. New York: Teachers College Press.

Eckert, Penelope. 2008. Variation and the indexical field. *Journal of Sociolinguistics.* 12:4. 453–76. doi:https://doi.org/10.1111/j.1467-9841.2008.00374.x.

Fairchild, Sarah, Ariel Mathis & Anna Papafragou. 2020. Pragmatics and social meaning: Understanding under-informativeness in native and non-native speakers. *Cognition* 200. 104171. doi:10.1016/j.cognition.2019.104171.

Fairchild, Sarah & Anna Papafragou. 2018. Sins of omission are more likely to be forgiven in non-native speakers. *Cognition* 181. 80–92. doi:10.1016/j.cognition.2018.08.010.

Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998–998. doi:10.1126/science.1218633.

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press. doi:10.1093/acprof:oso/9780198237907.001.0001.

Glass, Lelia. 2015. Strong necessity modals: Four socio-pragmatic corpus studies. *University of Pennsylvania Working Papers in Linguistics* 21(2).

Goodman, Noah D. & Michael C. Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences* 20(11). 818 – 829. doi:https://doi.org/10.1016/j.tics.2016.08.005. http://www.sciencedirect.com/science/article/pii/S136466131630122X.

Goodman, Noah D. & Andreas Stuhlmüller. 2013. Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science* 5(1). 173–184. doi:10.1111/tops.12007.

Hay, Jen, Paul Warren & Katie Drager. 2006. Factors influencing speech perception in the context of a merger-in-process. *Journal of Phonetics* 34. 458–484. doi:10.1016/j.wocn.2005.10.001.

Hay, Jennifer. 2009. Stuffed toys and speech perception. *Linguistics* 11(3). 135–163. doi:10.1515/LING.2010.027.

Henderson, R. & E. McCready. 2019. Signaling without saying: The semantics and pragmatics of dogwhistles. In preparation for Oxford University Press.

Huang, Y., E. Spelke & J. Snedeker. 2013. What exactly do number words mean? *Language Learning and Development* 9(2). 105–129. doi:https://doi.org/10.1080/15475441.2012.658731.

Hunt, Matthew & Eric K. Acton. 2022. "How's the wife?": Pragmatic reasoning in spousal reference. *Journal of Pragmatics* 188. 152–170. doi:10.1016/j.pragma.2021.11.005.

Irvine, Judith & Susan. Gal. 2000. Language ideology and linguistic differentiation. In *Regimes of language*, 35–83. Santa Fe, NM: School of American Research Press.

Irvine, Judith. 2001. S. 2001. Style as distinctiveness: the culture and ideology of linguistic differentiation. In John Rickford & Penelope Eckert (eds.), *Stylistic Variation,*, Cambridge University Press.

Jeong, Sunwoo. 2021. Deriving politeness from an extended lewisian model: The case of rising declaratives. *Journal of Pragmatics* 177. 183–207. doi:10.1016/j.pragma.2021.02.017.

Jones, Taylor, Jessica Kalbfeld, Ryan Hancock & Robin Clark. 2019. Testifying while black: An experimental study of court reporter accuracy in transcription of african american english. *Language* doi:10.1353/lan.0.0235.

Kao, Justine T., Jean Y. Wu, Leon Bergen & Noah D. Goodman. 2014. Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences* 111(33). 12002–12007. doi:10.1073/pnas.1407479111.

Karawani, Hadil & Brandon Waldon. 2017. Strength, duality and social meaning. *Amsterdam Colloquium* 21.

Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1)(1). 1–45. doi:10.1007/s10988-006-9008-0.

Kiesling, Scott F. 2016. Stance. In Kira Hall & Rusty Barrett (eds.), *Oxford Handbook of Language and Sexuality*, Oxford/New York: Oxford University Press.

King, Sharese. 2021. Rethinking race and place: The role of persona in sound change reversal. *Journal of Sociolinguistics* 25(2). 159–178. doi:10.1111/josl.12454.

Klecha, Peter. 2014. *Scalarity and Modality*: University of Chicago PhD dissertation.

Krifka, Manfred. 2007. Approximate interpretations of number words: A case for strategic communication. In Gerlof Bouma, Irene Krämer & Joost Zwarts (eds.), *Cognitive foundations of interpretation*, 111–126. Amsterdam: Royal Netherlands Academy of Arts and Sciences. doi:10.18452/9508.

Lasersohn, Peter. 1999. Pragmatic halos. *Language* 75(3)(3). 522–551. doi:https://doi.org/10.2307/417059.

Lassiter, Daniel & Noah D. Goodman. 2017. Adjectival vagueness in a bayesian model of interpretation. *Synthese* 194(10). 3801–3836. doi:10.1007/s11229-

015-0786-1.

Leffel, Timothy, Ming Xiang & Christopher Kennedy. 2016. Imprecision is pragmatic: Evidence from referential processing. *Semantics and Linguistic Theory* 26. doi:https://doi.org/10.3765/salt.v26i0.3937.

Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8. 339–359.

Mahler, Taylor. 2020. The social component of the projection behavior of clausal complement contents. *Linguistic Society of America* 5(1). 777. doi:10.3765/plsa.v5i1.4703.

Mason, Malia, Alice Lee, Elizabeth Wiley & Daniel Ames. 2013. Precise offers are potent anchors: Conciliatory counteroffers and attributions of knowledge in negotiations. *Journal of Experimental Social Psychology* 49. 759–763. doi:10.1016/j.jesp.2013.02.012.

Mazzarella, Diana, Emmanuel Trouche, Hugo Mercier & Ira Noveck. 2018. Believing what you're told: Politeness and scalar inferences. *Frontiers in Psychology* 9. 908. doi:10.3389/fpsyg.2018.00908.

Morzycki, Marcin. 2011. The several faces of adnominal degree modification. *Proceedings of West Coast Conference on Formal Linguistics 29*.

Niedzielski, Nancy. 1999. The effect of social information on the perception of sociolinguistic variables. *Journal of Social Psychology (Special Edition)* 18:1. 62–85. doi:https://doi.org/10.1177/0261927X99018001005.

Noveck, Ira. 2001. When children are more logical than adults: experimental investigations of scalar implicature. *Cognition* 78(2). 165 – 188. doi:https://doi.org/10.1016/S0010-0277(00)00114-1.

Papafragou, Anna & Julien Musolino. 2003. Scalar implicatures: experiments at the semantics–pragmatics interface. *Cognition* 86(3). 253 – 282. doi:https://doi.org/10.1016/s0010-0277(02)00179-8.

Pinkal, Manfred. 1995. *Logic and lexicon*. Dordrecht.

Podesva, Robert J. 2011. Salience and the social meaning of declarative contours: Three case studies of gay professionals. *Journal of English Linguistics* 39(3). 233–264.

Sauerland, Uli & Penka Stateva. 2011. Two types of vagueness. In *Vagueness and Language Use*, 121–145. London: Palgrave Macmillan. doi:10.1057/9780230299313_6.

Schwarz, Florian & Jeremy Zehr. 2021. Tutorial: Introduction to PCIbex – An Open-Science Platform for Online Experiments: Design, Data-Collection and Code-Sharing. *Proceedings of the Annual Meeting of the Cognitive Science Society* 43(43). https://escholarship.org/uc/item/1ng1q4c6.

Scontras, Gregory & Lisa Pearl. 2020. When pragmatics matters more for truth-value judgments: An investigation of quantifier scope ambiguity. *Unpublished* Re-

trived from https://www.socsci.uci.edu/~lpearl/papers/ScontrasPearl2021Manu_ChildScope.pdf.

Sikos, Les, Minjae Kim & Daniel J. Grodner. 2019. Social context modulates tolerance for pragmatic violations in binary but not graded judgments. *Frontiers in Psychology* 10. doi:10.3389/fpsyg.2019.00510.

Solt, Stephanie. 2014. An alternative theory of imprecision. *Semantics and Linguistic Theory* 24. doi:https://doi.org/10.3765/salt.v24i0.2446.

Solt, Stephanie, Chris Cummins & Marijan Palmovic. 2017. The preference for approximation. *International Review of Pragmatics* 9(2). 248 – 268. doi:https://doi.org/10.1163/18773109-00901010.

Staum Casasanto, Laura. 2008. *Experimental Investigations of Sociolinguistic Knowledge.*: Stanford University PhD dissertation.

Strand, Elizabeth A. 1999. Uncovering the role of gender stereotypes in speech perception. *Journal of Language and Social Psychology* 18(1). 86–100. doi:10.1177/0261927x99018001006.

Sumner, Meghan, Seung Kyung Kim, Ed King & Kevin B. McGowan. 2014. The socially weighted encoding of spoken words: a dual-route approach to speech perception. *Frontiers in Psychology* 4. doi:10.3389/fpsyg.2013.01015.

Syrett, K., C. Kennedy & J. Lidz. 2009. Meaning and context in children's understanding of gradable adjectives. *Journal of Semantics* 27(1). 1–35. doi:10.1093/jos/ffp011.

Syrett, Kristen & Athulya Aravind. 2021. Context sensitivity and the semantics of count nouns in the evaluation of partial objects by children and adults. *Journal of Child Language* 1–27. doi:10.1017/s0305000921000027.

Thomas, William C. 2021. Social inferences from the use of just as an exclusive particle. *Linguistic Society of America* 6(1). 746. doi:10.3765/plsa.v6i1.5009.

Van Der Henst, Jean-Baptiste, Laure Carles & Dan Sperber. 2002. Truthfulness and relevance in telling the time. *Mind and Language* 17. 457–466. doi:https://doi.org/10.1111/1468-0017.00207.

Wade, Lacey. 2022. Experimental evidence for expectation-driven linguistic convergence. *Language* doi:10.1353/lan.0.0257.

Waldon, Brandon & Judith Degen. 2020. Modeling behavior in truth value judgment task experiments. *Proceedings of the Society for Computation in Linguistics* 3(3). doi:10.7275/SG32-AQ30.

Welsh, Matthew, D Navarro & S. Begg. 2011. Number preference, precision and implicit confidence. In *Proceedings of the Annual Meeting of the Cognitive Science Society 33*, 1521–1526.

Xie, Guang-Xin & Ann Kronrod. 2012. Is the devil in the details? *Journal of Advertising* 41(4). 103–117. doi:10.1080/00913367.2012.10672460.

Yoon, Erica J., M. H. Tessler, Noah D. Goodman & Michael C. Frank. 2016. Talking

with tact : Polite language as a balance between kindness and informativity, .
Yoon, Erica J., Michael Henry Tessler, Noah D. Goodman & Michael C. Frank. 2020. Polite speech emerges from competing social goals. *Open Mind* 4. 71–87. doi:10.1162/opmi_a_00035.
Zhang, Qing. 2005. A Chinese Yuppie in Beijing: Phonological variation and the construction of a new professional identity. *Language in Society* 34. 431–66.
Zhang, Y. Charles & Norbert Schwarz. 2011. How and Why 1 Year Differs from 365 Days: A Conversational Logic Analysis of Inferences from the Granularity of Quantitative Expressions. *Journal of Consumer Research* 39(2). 248–259. doi:10.1086/662612. https://doi.org/10.1086/662612.

Andrea Beltrama
3401-C Walnut Street
Suite 300, C Wing
University of Pennsylvania
Philadelphia, PA 19104-6228
beltrama@sas.upenn.edu

Florian Schwarz
3401-C Walnut Street
Suite 300, C Wing
University of Pennsylvania
Philadelphia, PA 19104-6228
florians@sas.sas.upenn.edu

...