

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

From social identity to meaning interpretation: when looser speakers are treated more strictly

Permalink

<https://escholarship.org/uc/item/8wn9x1p8>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Beltrama, Andrea

Schwarz, Florian

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

From identity to interpretation: when looser speakers are treated more strictly

Andrea Beltrama (beltrama@sas.upenn.edu)

Department of Linguistics, 3401-C Walnut Street
Philadelphia, PA 19104 USA

Florian Schwarz (florians@ling.upenn.edu)

Department of Linguistics, 3401-C Walnut Street
Philadelphia, PA 19104 USA

Abstract

We explore the impact of speaker identity on the interpretation of number words in a T(ruth)-V(alue) J(udgment) task – a paradigm in which respondents assess whether a given description appropriately represents a given body of facts. 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, and especially so when participants do not identify with the speaker’s properties. This shows that TVJ assessments are impacted by respondents’ social perception of the speaker; but that they are affected by social considerations in a different way from other experimental tasks similarly tapping into meaning interpretation, suggesting a nuanced interplay between social information and pragmatic reasoning.

Keywords: pragmatics; social meaning; truth-value judgment-tasks; interpretation; personae

Introduction

Work in psycho- and sociolinguistics demonstrated that comprehenders track social information about the interlocutor when parsing linguistic input. Much of this work focused on how phonetic processing is affected by macro-sociological attributes of the speaker, including their location of origin, gender, or race (Niedzielski 1999; Babel 2012; Drager 2015; Sumner 2014; Wade 2022, i.a.), as well as by information about more specific “types” of people, or *personae*, salient in discourse (see D’Onofrio 2020 for an overview).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 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).

Much less investigated, however, is the question of how comprehenders’ perception of the speaker’s identity interacts with how they compute the meaning of a linguistic utterance. While research in psycholinguistics extensively explored how comprehenders reason about contextual information to interpret linguistic expressions (among many others: Tanenhaus

et al. 1995; Heller, Grodner & Tanenhaus 2008; Bott, Bailey & Grodner 2012; Schwarz 2019), this work by-and-large operated within gricean and neo-gricean frameworks, in which interlocutors are treated as homogeneous rational agents, and their distinctive social profiles is excluded from the array of factors that are normally taken to shape meaning computation and pragmatic reasoning – e.g., speakers’ intentions; conversational maxims; previous discourse; world knowledge (Grice 1975; Roberts 2012). However, the separation between social considerations and meaning interpretation has been called into question by work that highlighted the relevance of social factors to pragmatic reasoning, with an emphasis on politeness considerations. For example, quantifiers such as *some* are interpreted as lacking an upper-boundary – i.e., as meaning “some and possibly all” – in contexts in which listeners assume that the speaker is using them to be tactful towards the listener – e.g., in an utterance like “Some people hated your speech” (Bonnefon and Villejoubert 2006; Bonnefon et al. 2009; Mazzarella et al. 2018 i.a.).² In a similar vein, other social dimensions – e.g., speakers’ political orientation (Mahler 2020) and linguistic nativeness (Fairchild & Papafragou 2018) – have been suggested to affect meaning resolution. Taken together, and coupled with the insights from the sociolinguistic literature, these findings call for a more systematic investigation of how comprehenders reason about the speaker’s identity when interpreting meaning, suggesting that there would be much to gain from a deeper understanding of this process.

In this paper, we address this issue by asking how comprehenders’ perception of the speaker persona affects the way in which they assess the aptness of a description to represent a particular body of facts. Assessments of this sort are central to semantic interpretation, whose ultimate goal is to verify whether a given description correctly represents a particular state of affairs; in fact, they represent the core behavioral measure adopted in Truth Value Judgment (henceforth, **TVJ**) tasks (Crain & McKee 1985) – a paradigm extensively used to track the outcome of comprehenders’ interpretation in studies of meaning processing (Noveck, 2001; Papafragou & Musolino 2003; Bott & Noveck, 2004). For example, comprehenders’ acceptance of sentences such as “some giraffes

¹The 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.

²See also Brown and Levinson (1987); Terkourafi (2021) for theoretical implementations of politeness in pragmatic models.

have long necks” (Noveck 2001) has been taken to reveal that they interpreted quantifier *some* as lacking an upper-boundary (i.e. “some, and possibly all”). While recent work began to explore how respondents reason about the discourse situation and their relation to it to make these assessments (Sikos, Kim and Grodner 2019; Waldon & Degen 2020; Scontras & Pearl 2021), still largely uncharted is the role of social information in determining their outcome.

As a window into this question, we focus on (*im*)precision resolution, a processing challenge comprehenders need to navigate when parsing numerical and quantity expressions. For instance, an utterance such as “the price is \$200” can be assigned different interpretations – ranging from a maximally strict one (“exactly \$200”) to looser, more inclusive ones (e.g., “\$195-205”, “\$190-210”) – depending on the threshold of precision that interlocutors deem to be appropriate in the context (Lasersohn 1999)³. Imprecision resolution emerges as an especially suitable testbed for our purposes: it requires comprehenders to reason about different elements of the discourse context to arrive at an interpretation (Van der Henst et al. 2002; Cummins et al. 2012; Solt et al. 2017; Aparicio 2017); and this reasoning has recently been shown to be impacted by information about the speaker identity as well. In particular, Beltrama & Schwarz (2021, 2022) showed that respondents were more inclined to accept a slightly mismatching number as the referent of a numeral when the speaker embodied a “Chill” persona, socially expected to speak loosely, as opposed to a “Nerdy” persona, socially expected to speak more precisely; and that this effect was particularly pronounced for respondents who did *not* identify with the speaker persona. Evidence supporting these claims was based on participants’ *indirect* judgments about whether a certain fact could retroactively be accommodated as the referent of a numeral. Specifically, in a picture selection paradigm called the COVERED SCREEN TASK⁴, participants were presented with a numeral utterance (e.g., “The cost is \$200”) and a picture of a screen showing a slightly mismatching number (“\$212”; **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.

These findings raise an important question: does the perception of speaker persona similarly affect numeral interpretation when comprehenders are *directly* asked to determine the descriptive appropriateness of an utterance, as in a TVJ task? Besides allowing for a cross-paradigm validation of Beltrama & Schwarz findings and a deeper understanding of

how comprehenders reason about the speaker persona when interpreting an utterance, addressing this question would provide the opportunity of exploring TVJs’ sensitivity to social information – shedding light on a broader methodological issue central to experimental studies on meaning. To this end, we repurpose Beltrama & Schwarz’s (2022) original setup to a TVJ paradigm. If participants’ assessments are shaped by these factors in a parallel way to the judgments collected in the original study, we hypothesize that imprecise descriptions produced by Nerdy speakers should be interpreted more strictly, and hence accepted as appropriate *less* often than imprecise descriptions uttered by Chill speakers, especially for participants who do not identify with the speaker persona.

Methods

Participants 196 participants were recruited on Prolific and paid \$2 for participation. The study was implemented and administered online on the PC Ibex platform (see Schwarz & Zehr 2021 for details).⁵

Materials and procedure Following Beltrama & Schwarz (2021, 2022), we created dialogues in which one character asked a question and the other provided a numeral utterance in response, after checking their phone. Following the dialogue, participants were shown the image of the phone screen showing a number, and were told that it was the phone the speaker was looking at in the picture. Two factors were crossed in a 2x3 design: **Speaker Persona** and **Match**. Speaker Persona, implemented between subjects, manipulated the persona embodied by the speaker, with levels *Nerdy* (expected to speak precisely) and *Chill* (expected to speak more loosely). The social perception of these characters, as well as the expectations that they would elicit with respect to precision, was independently normed in a previous norming study (see Beltrama & Schwarz 2022 for details).

“Rachel and Arthur are looking for a one-way plane ticket”

“Alex and Eva are looking for a one-way plane ticket”



Figure 1: Nerdy characters

Figure 2: Chill characters

The Match factor manipulated how closely the uttered numeral (e.g., \$200, see Fig. 1-2) and the number on the phone

³See Lewis (1979), Pinkal (1995), Lasersohn (1999), Krifka (2007), Syrett (2009), Solt (2014), Klecha (2018), Beltrama & Hanink (2019), Beltrama (2021) for further work discussing imprecision across a variety of linguistic phenomena

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

⁵<https://farm.pcibex.net>



Figure 3: Filler item

matched, with 3 levels: *Match*, in which the displayed number is the same (\$200.00); *Mismatch*, in which the displayed number widely differs (\$650.12); or *Imprecise*, in which the displayed number differs only slightly (\$212.12).

24 items were counterbalanced across 4 lists, each with 6 items in *Match* and *Mismatch*, and 12 in *Imprecise*. In the *Imprecise* condition, the divergence between the uttered number and the one shown on the screen was comprised between 5% and 19%, with the range counterbalanced across items. The experimental items included numerals describing times (8), costs (8) and distances (8). 24 fillers were also included, in which participants saw characters embodying a third type of persona looking at their phone and producing utterances containing quantifiers *some* and *all*. The fillers were presented in alternation with the experimental items, such that participants never saw two occurrences of either type of item in a row. An example of a filler item is provided below.

At the end of each trial, participants were asked to assess whether, given the number shown on the phone, the utterance provided by the speaker provided an appropriate description of the current facts, with **Right** or **Wrong** as possible choices (**TVJ Question**). Figure 4 shows an example of a full trial.

Rachel and Arthur are looking for a one-way plane ticket



Here's the phone Arthur is looking at!



Is his response **right** or **wrong**?

F J

Figure 4: Full trial (Persona: Nerdy; Match: Imprecise)

An exit survey asked participants to answer to two further questions about the character on a 1 to 10 Likert scale.

- **Precision Expectation Question:** How precisely do you expect the character to speak?

- **Similarity Question:** To what extent do you see yourself in the stereotypical traits of these characters?

The Precision Expectation Question was included to confirm the assumption that Nerdy speakers should be expected to speak more precisely than Chill speakers. The Similarity Question was included to obtain an independent variable enabling us to test how the effect of Persona on participants' binary choices is modulated by comprehenders' own orientation towards these personae (see **H2** below).

Hypotheses

If participants' behavior in the TVJ task is affected by the experimental manipulations in a parallel way to the picture selection judgments from Beltrama & Schwarz (2021) study, we hypothesize the following:

- **H1:** Imprecise descriptions produced by Nerdy speakers should be interpreted more strictly than those uttered by Chill speakers, leading to higher rates of rejections in the *Imprecise* condition for Nerdy than Chill speakers.
- **H2:** The persona effect hypothesized in H1 should be especially prominent for participants who do not identify with the speaker persona – i.e., we should observe an inverse correlation between the difference in rejections rates associated with Chill vs. Nerdy speakers and the rating provided in response to the Similarity Question.

Results

As a first step, we verified the basic assumption linking participants' social personae and expectations of precision. Nerdy speakers were ascribed higher rates than Chill ones in the **Precision Expectation Question** (Nerdy: $M=5.24$; $sd=2.30$; Chill: $M=4.20$; $sd=2.03$); the difference was confirmed to be significant in a two-tailed t-test ($t(187)=3.24$; $p < 0.001$), suggesting that the assumption was met.

We then proceeded to consider participants' responses to the **TVJ Question**. Figure 5 shows the proportion of rejections (utterances labeled as "wrong") by Match and Persona, the two manipulations in the experiment, suggesting the following descriptive pattern: the overall rate of rejections is at floor and ceiling for the Match (1%) and MisMatch (99%) conditions respectively, and patterns in between for Imprecise (59%); furthermore, the proportion of rejections in the *Imprecise* condition is higher for Chill (65%) speakers than for Nerdy (54%) ones.

A mixed-effects logistic regression with **Match** and **Persona** (ref: *Chill*; *Imprecise*) as fixed effects and random intercepts for items and participants support these observations. Across personae, rates of rejections in the *Imprecise* condition were significantly higher than those in the *Match* condition ($\beta = 10.05$; $SE = 0.73$; $p < 0.001$) and lower than those in the *MisMatch* condition ($\beta = -6.74$; $SE = 0.46$; $p < .001$).⁶

⁶The analysis was carried out with the *lmerTest* package in R. Coefficients and p values were obtained via the "emmeans" R package.

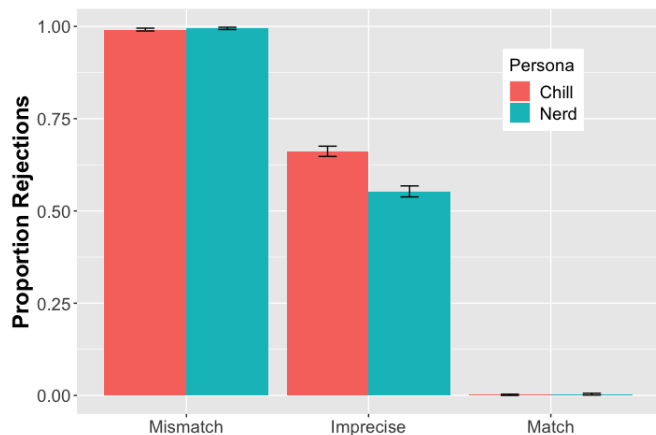


Figure 5: Rejections by Match and Persona

Moreover, a simple effect of Persona was found in the *Imprecise* condition, with rejections higher for Chill than for Nerdy speakers ($\beta = 1.03$; $SE = 0.47$; $p < .05$); no Persona effect was found for the *MisMatch* ($\beta = 0.91$; $SE = 0.98$; $p = 0.35$) or the *Match* condition ($\beta = 0.10$; $SE = 1.39$; $p = 0.93$).

To zero in on how the Persona effect is modulated by participants' own identification with the characters, we now focus on data from the *Imprecise* condition only. Figure 6 shows participants' choices by their self-ascribed degree of similarity to the speaker, corresponding to their response to the **Similarity Question**. The plot suggests that Chill speakers received higher rates of rejections at a lower degrees of similarity, while Nerdy speakers receive consistent responses across degrees of similarity.

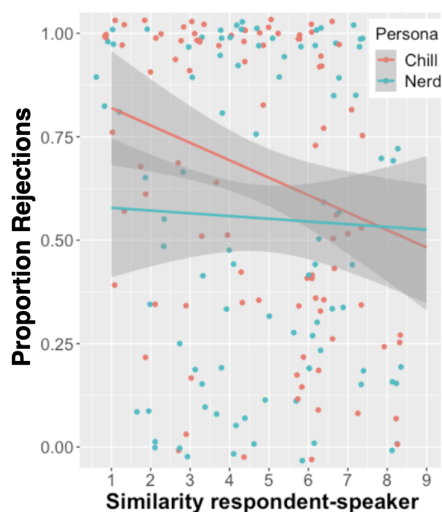


Figure 6: rejections by Persona and respondent-speaker similarity (Imprecise condition only)

We tested the pattern with a second ME regression on TVJ choices in the *Imprecise* condition data, with Persona (ref =

Chill) and Similarity (continuous, ref = 1) as predictors, and random intercepts for items and participants. We found an effect of Persona at low levels of participant-speaker similarity, with a higher rate of rejections for Chill speakers ($\beta = 3.17$; $SE = 1.29$; $p < 0.05$); and a simple effect of Similarity, with rejections for Chill speakers decreasing as similarity increases ($\beta = 0.45$; $SE = 0.17$; $p < 0.001$). Responses for Nerdy speakers were not affected by Similarity ($\beta = 0.03$; $SE = 0.17$; $p = 0.81$). As a result of rejections for Chill speakers decreasing as similarity increases, the difference between Chill and Nerd is neutralized at high degrees of Similarity, with the interaction between Similarity and Persona approaching significance ($\beta = 0.42$; $SE = 0.22$; $p = 0.08$).

Discussion

Two patterns emerge from our findings. First, comprehenders are *more* strict in judging descriptions uttered by speakers that they expect to speak *less* precisely, as suggested by the higher rate of rejections for Chill speakers than for Nerdy ones. This result suggests the *reverse* pattern of **Hypothesis 1**, which predicted that Nerdy speakers' utterances should have been judged more strictly than Chill speakers' ones. Second, respondents are especially uncharitable towards Chill speakers when they perceive themselves as dissimilar from them, as suggested by the higher rate of rejections attributed to Chill speakers by participants who do not identify with the speaker persona. This provides partial support to **Hypothesis 2**, which predicted that the Persona effect should be especially prominent for comprehenders who see themselves as different from the speaker: an effect of participants' self-identification with the speaker is indeed observed on responses to Chill speakers, but not to Nerdy ones.

Taken together, these results suggest that considerations about the persona embodied by the speaker and respondents' orientation towards it shape the outcome of the Truth Value Judgment task: a widely used paradigm in the study of meaning, whose sensitivity to comprehenders' social perception of the speaker remains largely unexplored. We take these observations to carry important theoretical and methodological implications for the study of meaning and language processing. Before considering them, however, it is important to first address a puzzle raised by our findings: why is the effect of persona on comprehenders' behavior in our study the opposite of the one observed in Beltrama & Schwarz COVERED SCREEN setup, in which participants rejected imprecise descriptions made by Nerdy speakers more often?

We propose that these different patterns of results can be illuminated by considering the distinct epistemic implications linked to rejecting a description across the two paradigms. In Beltrama & Schwarz (2021, 2022) task, rejecting the imprecise display and selecting the **COVERED** option is inconsequential in terms of speaker evaluation: 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 respondent to implying that the speaker is at fault for untruthful behavior – that is, that they are using language improperly, and that they’re therefore violating the assumption of cooperativeness that normally underlies conversational exchanges (Grice 1975). 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 conversational norms, leading respondents to be less charitable, hence more punitive, towards them. At the same time, because Chill speakers’ descriptions are actually expected to be less precise, they remain (in principle) more likely to be accepted in the presence of a slight mismatch between language and facts in a task where rejection doesn’t entail an ascription of untruthfulness – as in Beltrama & Schwarz (2021, 2022). The same line of reasoning allows us to capture the observation that responses to Chill speakers’ utterances, but not to Nerdy ones’, were affected by participants’ degree of self-identification with the speaker. If higher rates of rejections attributed to Chill speakers are indeed driven by a negative perception of this persona as prone to violating truthfulness, one could predict that the endorsement of this stereotype should be especially strong for people who do not see themselves as directly targeted by it, leading to uncharitable behavior in the responses; and should be weaker, and possibly absent, for people who instead do see themselves as directly involved with it, leading to a more lenient response behavior on their part. This is indeed what we observe: rates of rejections for Chill decrease as the self-perceived similarity between the respondent and the speaker increases, with the difference between Chill and Nerdy speakers eventually disappearing for participants who maximally identify with the Chill persona. On this view, the observed pattern aligns with the more general observation that negative stereotypes about a particular social group are more likely to affect language processing for comprehenders who do not see themselves as part of that group, as suggested by findings from studies on phonetic processing (Niedzielski 1999; Wade 2022 i.a.).

On a broad level, these results enrich our understanding of how social information shapes language processing. As summarized in the Introduction, a growing number of studies in sociolinguistics showed that listeners’ speech perception of linguistic input, particularly at the phonetic level, is affected by speakers’ demographic features (e.g., Niedzielski 1999; Babel 2012; Drager 2015; Wade 2022), as well as by the social personae they embody in the discourse context (D’Onofrio 2015; 2018; 2020). Our findings broaden the scope of such endeavors by highlighting personae as a source of contextual information shaping meaning resolution and pragmatic reasoning – a domain of language processing that has been generally taken to be independent of interlocutors’ social profile. Specifically, we have shown that

comprehenders reason about the social type embodied by the speaker to 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 more specific ramifications for the study of meaning – both theoretically and methodologically.

On a theoretical level, our approach provides the opportunity of developing a novel, socially integrative perspective on meaning-related phenomena such as (im)precision and the use of quantity expressions. On this view, our findings connect with an emergent body of work at the intersection of semantics, pragmatics and sociolinguistics which demonstrated that comprehenders can infer social information about speakers by reasoning about linguistic forms’ semantic and pragmatic properties (Acton & Potts 2014; Acton 2019; Glass 2015; Beltrama 2018; Beltrama & Staum Casasanto 2022; Beltrama, Solt & Burnett 2022; Thomas 2021; Jeong 2021; Hunt & Acton 2022; see Beltrama 2020 for an overview). Combined with Beltrama & Schwarz (2021, 2022), our results crucially enrich this outlook by showing that interlocutors’ ability to compute the meaning conveyed by a linguistic form cannot be divorced from interlocutors’ ability to associate the use of this form with specific, socially recognizable categories of speakers. The emerging picture is one in which processes of person perception and meaning interpretation are bi-directional: social information about the speaker identity, including the persona that they embody, is inferrable from the semantic and pragmatic properties of an utterance, and conversely shapes how interlocutors compute the message conveyed by the utterance. It follows that phenomena such as (im)precision are best framed within a socially informed approach to semantics and pragmatics: one in which the cognitive processes whereby interlocutors compute the message conveyed by an utterance are crucially tied to the social context in which interaction takes place; and in which inferential processes at the social and propositional level are treated as (at least partially) overlapping (see Burnett 2019 for a formalization).

On a methodological level, 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 (see Introduction). 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, Kim and Grodner 2019; Scontras and Pearl 2021); the probability with which respondents see themselves as producing the utterance (Waldon and Degen 2020); or properties inherent to the de-

mographic background of the speaker, such as their linguistic nativeness (Fairchild and Papafragou 2018). Our results open a novel perspective on the context-sensitivity of this behavioral measure, suggesting two takeaways. First, when making these assessments, comprehenders engage in social reasoning about two related, yet distinct component of the discourse context: the social identity of the speaker; and their own orientation towards and this identity. Second, the way in which these factors impact comprehenders' determination of an utterance appropriateness in a TVJ task does not align with the way in which the same factors affect other measures similarly aimed at tapping into meaning interpretation. Especially noteworthy, here, is that imprecise interpretations of utterances produced by socially imprecise speakers are *accepted* more often in a task requiring participants to retroactively determine whether a fact fits a particular utterance, but are *rejected* more often when participants are asked to explicitly assess the appropriateness of an utterance, as in the current task. This indicates 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. While this interaction has been extensively investigated in philosophical approaches to meaning and conversation (see Fricker 2007 on *testimonial injustice*), it remains crucially underexplored in semantic and pragmatic ones, 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 who is asked to assess its viability.

References

- Acton, E. (2019). Pragmatics and the social life of the english definite article. *Language*, 95(1), 37–65. doi: 10.1353/lan.2019.0010
- Acton, E., & Potts, C. (2014). That straight talk. Sarah Palin and the sociolinguistics of demonstratives. *Journal of Sociolinguistics*, 18(1), 3–31. doi: <https://doi.org/10.1111/josl.12062>
- Aparicio, H. (2017). *Processing context-sensitive expressions: the case of gradable adjectives and numerals*. Unpublished doctoral dissertation, University of Chicago.
- Babel, M. (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, A. (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, A. (2020). Social meaning in semantics and pragmatics. *Language and Linguistics Compass*, 14(9). doi: 10.1111/Inc3.12398
- Beltrama, A. (2021). Just perfect, simply the best: an analysis of emphatic exclusion. *Linguistics and Philosophy*. doi: 10.1007/s10988-021-09326-x
- Beltrama, A., & Hanink, E. (2019). Marking imprecision, conveying surprise: *Like* between hedging and mirativity. , 55(1), 1–34. doi: 10.1017/S0022226718000385
- Beltrama, A., & Schwarz, F. (2021). Imprecision, personae, and pragmatic reasoning. *Semantics and Linguistic Theory*, 31, 122. doi: 10.3765/salt.v31i0.5107
- Beltrama, A., & Schwarz, F. (2022). *Imprecision and social identity: How persona-level information affects meaning resolution*. (Manuscript under review.)
- Beltrama, A., Solt, S., & Burnett, H. (2022). *Context, precision and social perception: a socio-pragmatic study*. (In press at *Language in Society*)
- Beltrama, A., & Staum Casasanto, L. (2022). The social meaning of semantic properties. In L. Hall-Lew, E. Moore, & R. J. Podesva (Eds.), *Social meaning and linguistic variation: Theorizing the third wave*. Cambridge University Press. doi: 10.1017/9781108578684.004
- Bonnefon, J.-F., Feeney, A., & Villejoubert, G. (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, J.-F., & Villejoubert, G. (2006). Tactful or doubtful?: Expectations of politeness explain the severity bias in the interpretation of probability phrases. *Psychological Science*, 17(9), 747–751. (PMID: 16984289) doi: 10.1111/j.1467-9280.2006.01776.x
- Bott, L., Bailey, T. M., & Grodner, D. (2012). Distinguishing speed from accuracy in scalar implicatures. *Journal of Memory and Language*, 66(1), 123 – 142. doi: 10.1016/j.jml.2011.09.005
- Bott, L., & Noveck, I. A. (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
- Brown, P., & Levinson, S. C. (1987). *Politeness*. Cambridge: Cambridge University Press.
- Burnett, H. (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>
- Coupland, N. (2007). *Style: Language variation and identity*. Cambridge University Press.
- Crain, S., & McKee, C. (1985). Acquisition of structural restrictions on anaphora. *North East Linguistics Society*, 16(1).
- D'Onofrio, A. (2018). Personae and phonetic detail in sociolinguistic signs. *Language in Society*, 47(4), 513–539. doi: 10.1017/S0047404518000581
- D'Onofrio, A. (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, A. (2020). Personae in sociolinguistic variation. *WIREs Cognitive Science*, 11(6). doi: 10.1002/wcs.1543
- Drager, K. K. (2015). *Linguistic variation, identity construction and cognition*. Language Science Press. doi: 10.17169/LANGSCI.B75.22
- Eckert, P. (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, S., & Papafragou, A. (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
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. doi: 10.1093/acprof:oso/9780198237907.001.0001
- Glass, L. (2015). Strong necessity modals: Four socio-pragmatic corpus studies. *University of Pennsylvania Working Papers in Linguistics*, 21(2).
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics, Vol. 3, Speech Acts* (p. 41–58). New York: Academic Press.
- Huang, Y., Spelke, E., & Snedeker, J. (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, M., & Acton, E. K. (2022). “How’s the wife?”: Pragmatic reasoning in spousal reference. *Journal of Pragmatics*, 188, 152–170. doi: 10.1016/j.pragma.2021.11.005
- Jeong, S. (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
- Kiesling, S. F. (2016). Stance. In K. Hall & R. Barrett (Eds.), *Oxford handbook of language and sexuality*. Oxford/New York: Oxford University Press.
- King, S. (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, P. (2018). On unidirectionality in precisification. *Linguistics and Philosophy*, 41(1), 87–124. doi: <https://doi.org/10.1007/s10988-017-9216-9>
- Krifka, M. (2007). Approximate interpretations of number words: A case for strategic communication. In G. Bouma, I. Krämer, & J. Zwarts (Eds.), *Cognitive foundations of interpretation* (pp. 111–126). Amsterdam: Royal Netherlands Academy of Arts and Sciences. doi: 10.18452/9508
- Lasersohn, P. (1999). Pragmatic halos. *Language*, 75(3)(3), 522–551. doi: <https://doi.org/10.2307/417059>
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- Mahler, T. (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
- Mazzarella, D., Trouche, E., Mercier, H., & Noveck, I. (2018). Believing what you’re told: Politeness and scalar inferences. *Frontiers in Psychology*, 9, 908. doi: 10.3389/fpsyg.2018.00908
- Niedzielski, N. (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, I. (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](https://doi.org/10.1016/S0010-0277(00)00114-1)
- Papafragou, A., & Musolino, J. (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](https://doi.org/10.1016/s0010-0277(02)00179-8)
- Pinkal, M. (1995). *Logic and lexicon*. Dordrecht.
- Podesva, R. J. (2011). Saliency and the social meaning of declarative contours: Three case studies of gay professionals. *Journal of English Linguistics*, 39(3), 233–264.
- Roberts, C. (1996/2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6), 1–69. doi: 10.3765/sp.5.6
- Schwarz, F. (2019). Presuppositions, projection, and accommodation. In C. Cummins & N. Katsos (Eds.), *The oxford handbook of experimental semantics and pragmatics*. Oxford University Press. doi: 10.1093/oxfordhb/9780198791768.001.0001
- Schwarz, F., & Zehr, J. (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). Retrieved 2021-11-10, from <https://escholarship.org/uc/item/1nqlq4c6>
- Sikos, L., Kim, M., & Grodner, D. J. (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, S., Cummins, C., & Palmovic, M. (2017). The preference for approximation. *International Review of Pragmatics*, 9(2), 248 – 268. doi: <https://doi.org/10.1163/18773109-00901010>
- Sumner, M., Kim, S. K., King, E., & McGowan, K. B. (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., Kennedy, C., & Lidz, J. (2009). Meaning and context in children’s understanding of gradable adjectives. *Journal of Semantics*, 27(1), 1–35. doi: 10.1093/jos/ffp011
- Tanenhaus, M., Spivey-Knowlton, M., Eberhard, K., & Sedivy, J. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217), 1632–1634. doi: 10.1126/science.7777863
- Thomas, W. 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, J.-B., Carles, L., & Sperber, D. (2002). Truthfulness and relevance in telling the time. *Mind and Language*, 17, 457–466. doi: <https://doi.org/10.1111/1468-0017.00207>
- Wade, L. (2022). Experimental evidence for expectation-driven linguistic convergence. *Language*. doi: 10.1353/lan.0.0257
- Waldon, B., & Degen, J. (2020). Modeling behavior in truth value judgment task experiments. *Proceedings of the Society for Computation in Linguistics*, 3(3). doi: 10.7275/SG32-AQ30