Social identity affects imprecision resolution – and in different ways for different tasks

Abstract. In addition to a descriptive meaning, linguistic utterances carry a *socio-indexical* meaning – i.e., information related to speakers' identity and personality features. While it's been suggested that these two dimensions might affect one another, little is known about how interlocutors jointly navigate them throughout the interpretation process. We address this issue by asking how comprehenders' interpretation of numerals is affected by the speaker's *social persona* – and in particular, by whether the speaker embodies a *Nerdy* persona, socially expected to describe things precisely, vs. a *Chill* one, expected to be less precise. Evidence from a picture selection task suggests that comprehenders base their interpretation of numerals on higher standards of precision when these are uttered by Nerdy characters (Exp1); but evidence from a Truth-Value Judgment task shows that conorehenders exhibit more tolerance for imprecision towards Nerdy speakers when judging the appropriateness of an imprecise description to represent a given fact (Exp2). We explain these findings by arguing that persona-based information can affect meaning interpretation in two ways: by shifting the standard of (im-)precision adopted in deciding what the circumstances described can be like based on socially driven expectations about the speaker; or by modulating the degree of charity extended to the speaker in assessing what they said as right or wrong. We take these findings to highlight the importance of incorporating the interplay of social and descriptive meaning into our understanding of pragmatic reasoning, and to reveal different sensitivities of minimally varied experimental methodologies to social considerations.

Word count: 7872

1 Introduction

Linguistic utterances don't just allow interlocutors to describe reality; they also carry a *socio-indexical* meaning, conveying information about who the speakers are – their demographic background, ideological orientation, and personality traits. The socio-indexical dimension of meaning has been at the center of the focus of work in sociolinguistics and linguistic anthropology, as part of the broader endeavor of unpacking the relationship between linguistic variation and the socio-cultural context in which humans operate (Labov 1966; Ochs 1992; Silverstein 2003; Agha 2003; Eckert 2008; Podesva 2011; Gal and Irvine 2019; Campbell-Kibler 2011; Levon 2016; see Hall-Lew et al. 2021 for a recent overview). Among many other contributions, this work highlighted two distinctive aspects of the way in which social meanings are conveyed in communication. First, the socio-indexical content of linguistic forms does not boil down to an enumeration of speakers'

demographic traits – e.g., their age, gender or geographical origin; rather, it includes a multi-dimensional package of different traits that bear on more specific aspects of the speaker's identity and personality, and tend to assemble in holistic, stereotypically salient *personae* (Irvine 2001; Agha 2005; Coupland 2007; Eckert 2000, 2008; Podesva 2011; D'Onofrio 2018; Gal and Irvine 2019). Examples of personae systematically invoked by particular clusters of linguistic features include "Jocks" and "Burnouts" (Eckert 2008), "Valley Girls" (D'Onofrio 2015), "Nerds" (Bucholtz 2001), "Frat Boys" (Kiesling 2018), among many others (see D'Onofrio 2020 for an overview). Second, social meanings are not just highly salient in interaction, but do impact language processing at the cognitive level, as extensively shown in connection to different domains of sound production and perception (Strand 1999; Niedzielski 1999; Hay 2009; Babel 2012; Drager 2015; Sumner et al. 2014; Wade 2022 i.a.). Again, such effects are not limited to the macro-social, demographic features of speakers, but are also found in connection to more specific persona-level constructs: for example, listeners have been shown to adjust their perception of vowel boundaries when the speaker embodies a "Valley Girl" persona, stereotypically associated with the state of California, revealing an impact of these constructs on speech categorization (D'Onofrio 2015). Similar effects have been shown in connection to other domains of language processing such as assessments of "foreign accented" speech (D'Onofrio 2019) and syntactic parsing (Choe et al. 2019).

While central to sociolinguistic research, the category of social meaning has received less attention in semantics and pragmatics. Of course, much work at the semantics/pragmatics interface revolves around the key idea that utterance interpretation requires comprehenders to engage in active inferential work to interpret linguistic expressions (Grice 1975; Horn 1984; Gazdar 1980; Levinson 2000; Roberts 2012 i.a.), and explored the cognitive processes giving rise to such inferences (Noveck 2001; Papafragou and Musolino 2003; Bott et al. 2012; Doran et al. 2012; Degen and Tanenhaus 2015; Huang and Snedeker 2018). However, the predominant perspective adopted in this work has been one in which the distinctive identity and personality features of interlocutors are only tangential to utterance interpretation. The roots of this trace back to Grice's foundational theory of communication, and in particular to its framing of conversational partners as idealized, socially undifferentiated rational agents who can be expected to resolve the meaning of an utterance in essentially the same way, regardless of who they are – an idea crucially reflected in the notion that conversational inferences must be *calculable* in a very general sense (Grice 1975).

In recent years, a growing body of research has enriched this perspective by extending the scope of the semantics/pragmatic interface to include dynamics pertaining to the social domain broadly construed. For example, it has been shown that pragmatic reasoning is shaped by social factors such as politeness (Bonnefon et al. 2009; Yoon et al. 2016, 2020; Mazzarella et al. 2018); affect (Kao et al. 2014; Bergen 2016); and speaker-specific information such as linguistic nativeness (Fairchild and Papafragou 2018) or political orientation (Henderson 2019; Mahler 2020, 2022). In a parallel vein, the development of models of communication such as the *Rational Speech Act* framework (Frank and Goodman 2012; Goodman and Stuhlmüller 2013; Lassiter and Goodman 2017; Goodman and Frank 2016 i.a.), which frame utterance interpretation as a goal-oriented, domain-general activity that is part of *social cognition* more broadly, has led to a broadening of the empirical and formal approaches to the study of pragmatic inferences. Drawing on these insights, proposals have been developed to capture the signaling and uptake of social meanings with formal tools similar to those deployed to formalize pragmatic inferences – an endeavor reflected in both probabilistic, game-theoretic approaches (Burnett 2017, 2019; Henderson and McCready 2019)

and neo-gricean maxim-based frameworks (Acton 2019, 2022). Finally, recent work looking at a variety of linguistic phenomena has shown that comprehenders promptly infer identity and personality features of a speaker from the semantic and pragmatic properties of their utterances (Acton and Potts 2014; Beltrama and Staum Casasanto 2017, 2022; Acton 2019; Glass 2015; Karawani and Waldon 2017; Jeong 2021; Thomas 2021; Hunt and Acton 2022; see Beltrama 2020 for an overview). A takeaway of this research is that comprehenders closely track different aspects of the descriptive dimension of meaning to infer socio-indexical information.

Yet, despite these insights, much work at the semantics/pragmatics has continued to operate under the assumption that the distinctive identity and personality features of interlocutors are peripheral, at best, to utterance interpretation. As a result, a striking disconnect persists between the wealth of socio-indexical information typically available to comprehenders in communication – who they are, what personae they embody, and how these features are socially perceived by an interlocutor – and the reasoning process that is normally seen as central component to meaning interpretation. As a consequence, much remains to be seen about the dynamics whereby comprehenders recruit social meaning to zero in on the descriptive content conveyed by an utterance; and how the relation between meaning interpretation and socio-indexical information should be captured in broader frameworks for pragmatic reasoning. We believe that addressing these questions is important not just in light of the growing evidence of an interaction between the descriptive and social dimensions of meaning reviewed above; but also in light of the more general endeavor of framing pragmatic reasoning within the study of human action and cognition. In fact, it has been shown that stereotypical and persona-based representations play a major role in many domains of behavior: in addition to sociolinguistic work on phonetic processing mentioned above, this includes research on phenomena beyond the domain of language proper, including decision making (see Fiske 2018; Jenkins et al. 2018; Stolier et al. 2020) and the calibration of the epistemic authority interlocutors are willing to ascribe to others. For example, work in philosophy has highlighted the role of social stereotypes as a driving force behind prejudicial behaviors linked to *credibility deficit* – e.g., when a woman is treated as lacking authority on a matter on account of particular stereotypes (Fricker 2007) – and *credibility excess* – e.g. when someone asks an Asian-American person seated nearby to help them with their math problem, relying on the stereotype that everyone embodying this social identity is proficient in the discipline (Davis 2016). It follows that, if pragmatic reasoning is to be seen as sharing a common denominator with these domains of cognition, it is crucial to better understand its sensitivity to the relevant types of social information – which have indeed been shown to be highly impactful for such other realms of human action.

In this paper, we take a first step in this direction by asking how persona-based social information affects two complementary aspects of pragmatic reasoning in evaluating utterances: how comprehenders infer details about a described state of affairs based on a given statement; and how comprehenders assess whether the description of a given fact that a particular statement provides is right or wrong. We explore these questions by looking at the phenomenon of (potentially) *imprecise* uses of numerals.

2 Imprecision and numerals: a case study

As is well known, speakers routinely use numerical expressions in a way that isn't fully adherent to the facts being described. For example, someone could reasonably utter (1) when the time is in

fact 6:03; or (2) when the actual price is $295.

(1)   It's 6 o'clock.

(2)   The ticket costs $300.

This phenomenon, traditionally known as *imprecision*, has been extensively investigated in philosophy, semantics and pragmatics, both from a formal (Austin 1962; Lewis 1979; Pinkal 1995; Lasersohn 1999; Kennedy 2007; Sauerland and Stateva 2007, 2011; Solt 2014; Klecha 2018) and an experimental perspective (Van Der Henst et al. 2002; Syrett et al. 2009; Cummins et al. 2012; Leffel et al. 2016; Aparicio et al. 2016; Aparicio 2017; Syrett and Aravind 2021). Two distinctive properties of (im)precision, in particular, make it an ideal testbed for our question.

One property is that the level of precision with which numerals are used has already been shown to serve as a productive cue for listeners to draw social inferences about speakers' identity and personality. In particular, speakers using sharp vs. round numbers – normally taken to signal a high vs. lower level of precision (Krifka 2007) respectively – are associated with distinct clusters of social qualities (Beltrama 2018; Beltrama et al. 2022). On the one hand, precise speakers are perceived as featuring high *intellectual status* – e.g., as articulate, educated, intelligent; and as having low *sociability* – e.g., as annoying, pedantic, obsessive, and generally unlikable.[1] On the other hand, 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.[2]

The second relevant property is that the possibility of speaking imprecisely requires comprehenders to perform inferential work in interpreting numerals. This can be observed in two separate, complementary types of communicative situations. First, when a comprehender is construing a representation of what facts are being described by an utterance, they have to determine the extension of the numeral in a given context – i.e., what range of values the numeral can be taken to represent: a description such as "The ticket costs $200", for instance, can reasonably be taken as describing price ranges of varying size – e.g., the exact price of $200; a narrow interval from $195 to $205; a broader interval comprised from $210 to $190, etc. To zero in on the relevant range of values being represented, a comprehender is therefore required to engage in reasoning about the specifics of the conversational setting – and thus ultimately settle on an appropriate level of precision, and on the extension of the numeral. For example, the description above is more likely to be interpreted as referring to a broad range of values in a context that doesn't make details particularly relevant (e.g., a quick internet search to get a sense of the price range of the tickets); and to refer to a narrower range, or even the exact number, in a context that makes details relevant (e.g., a purchase under a tight budget). But comprehenders may also have to engage in inferential work in another type of situation: one in which they have direct access to a given set of circumstances, and they may have to evaluate whether a particular numerical description deployed by a speaker should be treated as right vs. wrong in light of these facts. Here, the reasoning crucially involves the comprehender

---

[1]In the social psychology literature, these two dimensions are commonly referred to as *Competence*, or "intellectually good" and *Warmth*, or "socially good" (see Fiske 2018 for an overview). In the sociolinguistics literature, these two dimensions are commonly referred to as *Status* and *Solidarity* (see Milroy and Preston 1999; Campbell-Kibler 2010 for overviews). For the purposes of the present article, we see these labels as effectively interchangeable.

[2]See Welsh et al. (2011); Mason et al. (2013); Xie and Kronrod (2012); Zhang and Schwarz (2011); Pena-Marin and Bhargave (2016) for similar findings in the marketing and social psychology literature.

calibrating how much *leniency* they are willing to extend toward a speaker describing facts imprecisely. For example, given a known price of \$207, a comprehender may need to adjudicate whether a statement like "the price is \$200" can be taken to be a close enough representation of this price in the context, or should instead be rejected as a mischaracterization thereof – and thus react accordingly in a conversation.

Drawing on these two properties of numerals, we thus ask whether, and how, comprehenders utilize socio-indexical information about the speaker to engage in these domains of pragmatic reasoning. Our studies shed light on this question by directly comparing comprehenders' interpretation of numeral utterances produced by two speakers embodying stereotypical social personae associated with higher vs. lower precision standards: a *Nerdy* speaker, incarnating a constellation of social qualities associated with high levels of precision; and a *Chill* speaker, incarnating social qualities associated with lower levels of precision. In the next section, we turn to grounding the choice of this contrast in the sociolinguistics literature and spell out specific hypotheses with respect to the possible effects of persona-based interpretation on the interpretation of numerals.

3   The persona contrast: implementation and hypotheses

Even though descriptive precision has not been linked to specific stereotypes in the sociolinguistics literature, prior work on social meaning does provide viable starting points for hypothesizing a connection between variation in precision and different social personae. In particular, the persona of a Nerdy character has been explicitly linked to linguistic manifestations of detail-orientedness, such as *hyper-articulation* (Bucholtz 1999, 2001). This style of utterance production is characterized by phonetic forms displaying a high degree of detail along one or more dimensions (e.g., expanded vowel length, enhanced intensity on consonant release, reduced co-articulation; see Lindblom 1990; de Jong et al. 1993), and has been argued to convey individual social qualities very similar to those evoked by precise numerals, including *articulateness* (Podesva et al. 2015), *learnedness* (Bucholtz 2001; Benor 2004), *effortfulness* (Eckert 2008), and *detail-orientedness* (Podesva 2007). By the same token, stereotypical personae such as "Surfer dudes", "Skaters", and "Frat Boys" have been linked to the cluster of qualities such as laid-backness, chillness and nonchalance (see Kiesling 2004, 2018), all of which are also central, or at least closely related, to the socio-indexical profile of imprecise numerals. Against this background, the opposition between Nerdy vs. Chill speakers – once properly implemented and adequately normed – seems like a promising candidate for embodying the contrast between the two clusters of qualities that have been shown to be associated with high vs. low precision: high intellectual standing, learnedness and educatedness, pedantry and purposeful un-coolness for the former; laid-backness, friendliness and sociability, and a low investment in projecting intellectual stature for the latter. On the basis of these considerations, we utilize these two personae to test how the interpretation of numeral expressions is affected by the distinctive social profile of the speaker. Representations of these personae were implemented visually using cartoon images of two sets of characters engaged in conversation:[3] one cartoon involved two stereotypically nerdy characters, called Arthur and Rachel; the other involved two chill characters, named Alex and Eva. The two sets of characters are displayed

---

[3]The cartoons were drawn using the cartoon drawing software Pixton: `https://www.pixton.com`.

in Figure 1-2.[4]

Figure 1: Nerdy characters

Figure 2: Chill characters

Our general hypothesis is that, everything else being equal, Nerdy speakers should be expected to describe things more precisely than Chill speakers; and that this expectation should impact interpretation across different domains of reasoning, such as those discussed above. Concerning the first – the determination of the range of values within the extension of the numeral – numerals uttered by Nerdy speakers should then be interpreted more precisely than numerals uttered by Chill ones, and thus associated with a narrower range of values on average. We will refer to this as **Hypothesis 1**.

With regards to the second type of situation – i.e., the adjudication of imprecise statements vis-a-vis a known fact –, there are two plausible outcomes one could imagine: one is that speaker persona might impact comprehenders' assessment of the acceptability of these statements in a way parallel to the above. That is, imprecise descriptions produced by Nerdy speakers, by virtue of being associated with a narrower range of values, should be rejected as mis-characterizations *more* often than imprecise descriptions produced by Chill speakers. We call this **Hypothesis 2A**. An alternative possibility is that comprehenders might determine whether to accept or reject an imprecise statement by recruiting social information for reasoning on an epistemic level – namely, to assess the amount of *credibility* that they are willing to ascribe to a speaker producing an imprecise statement. Accordingly, since Nerdy speakers are generally (perceived to be) more accurate in their descriptive endeavors, their utterances could be viewed more charitably than those by Chill ones, due to their (perceived) greater likely for have some valid reason to speak imprecise in the specific context; by contrast, Chill ones, who tend to be more approximate, may be seen as less worthy of being given the benefit of the doubt, and thus treated less charitably at the interpretive level. This should lead to the opposite pattern of the one above, which we refer to as **Hypothesis 2B**:

---

[4]Note that there are many interesting questions and further possibilities arising from other potential combinations of the personae and gender of the speaker and hearer, which should be explored in future work. Here, we opted for having both characters in each type of conversation represent the same persona as a simple first test case, and also kept the relation between gender and interlocutor role in conversation constant.

imprecise descriptions produced by Nerdy speakers should be rejected as mis-characterizations *less* often, and instead be accepted as close-enough to being true more often – than imprecise descriptions produced by Chill speakers.

To test these hypotheses, we proceed in two steps. First, we norm our implementation of the Nerdy and Chill personae (§4); second, we deploy two minimally different experimental paradigms to tap into how comprehenders recruit these social personae when reasoning about the descriptive content of numerals along the two dimensions discussed above (§5-6).

4   Norming study: Establishing the link between Persona and (im)precision

Our first step involves ascertaining the hypothesized association between the two sets of Nerdy vs. Chill characters illustrated above and precision in speech. To this end, we conducted a norming study: 240 participants (Median Age = 29; female = 171; male = 67; other = 2) recruited on Prolific were shown the pictures in Figure 1 in a between-subject design, with half seeing the Nerdy characters, and the other half the Chill ones. Participants had to perform two tasks. First, they were asked to list three attributes and a stereotypical label for the characters shown, to confirm the contrast between the two sets of characters with regards to their social properties. Participants' responses are illustrated in the word clouds in Figures 3-4, and confirm that the perception of the Nerdy vs. Chill characters largely aligns with what we aimed for: the former are overwhelmingly seen as embodying social qualities indicative of high intellectual standing (e.g., clever, smart) and introvert personality (e.g., quiet; awkward) and are consistently associated with a Nerdy/Geeky stereotype. In contrast, the latter are ascribed attributes such as chill, laid-back, relaxed, easy, and cool; and in addition, a sociable personality (e.g., friendly; outgoing).[5] While slightly less homogeneous, the stereotypical associations for the Chill characters also broadly align with a coherent set of personae linked to combinations of these qualities.
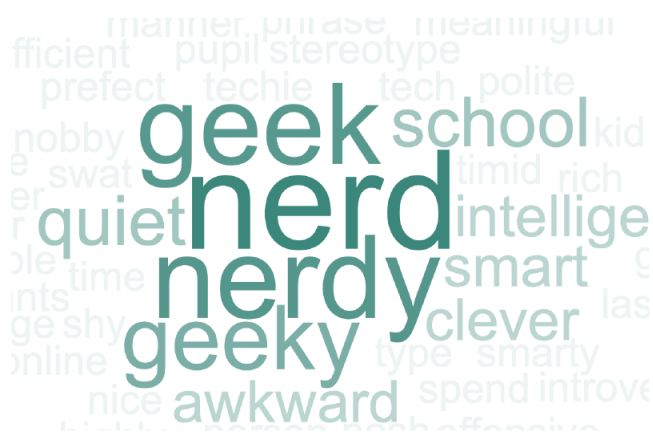

Figure 3: Qualities ascribed to Nerdy characters.   Figure 4: Qualities ascribed to Chill characters.

In the second task, participants were asked to rate how precisely they thought the characters

---

[5]The word clouds reported in Figures 3-4 were created via the Word Cloud visualization function in Qualtrics.

shown to them would speak about times and quantities on a 1-10 scale (10=maximally precise). This aimed to confirm the hypothesis that, in addition to being perceived as socially different, the two characters were also linked to different expectations with regard to the degree of precision in their speech. Figure 5 suggests that the average precision ratings for the Nerdy and the Chill characters indeed differed, the former being expected to speak more precisely about quantities than the latter (M= 7.20; sd = 2.36 vs. M = 5.63; sd = 2.29). This difference was confirmed to be significant by a paired two-tailed t-test ($t(238.67)= 5.23$; $p < 0.001$). Having established the viability of the persona contrast, we can proceed to test our hypotheses.
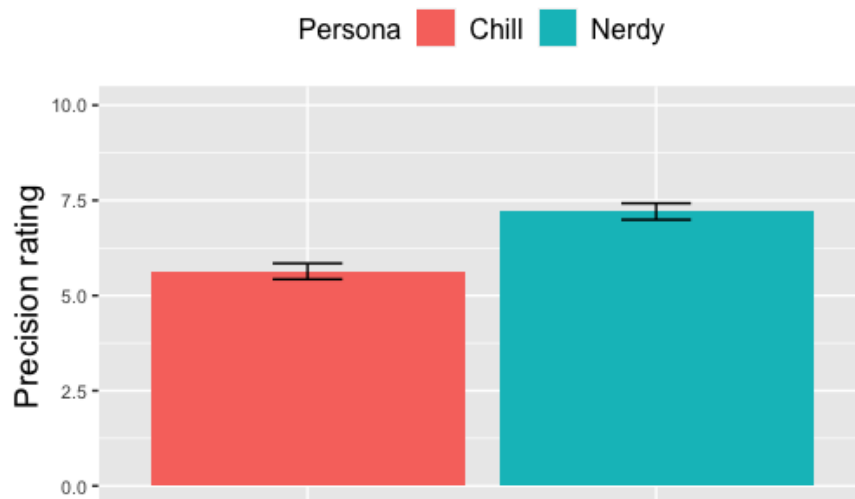
Figure 5: Expected precision of Chill vs. Nerdy character

5 Experiment 1: Inferring facts from statements

Experiment 1 explores how persona-based information affects comprehenders in determining the range of values corresponding to a numeral description in a given context. We utilize a paradigm that we call the COVERED SCREEN TASK, inspired by the *covered box* picture selection task from the experimental semantics literature (e.g., Huang et al. 2013; Schwarz et al. 2016).

5.1 Methods & Design

Our stimuli utilized visually displayed dialogues, like those in the norming study reported above. A set of characters was presented in a conversation, prefaced by a brief context sentence establishing the purpose of the conversation (e.g., looking for a plane ticket); in the dialogue, the female character would ask a question, and the male character would respond based on information he accessed by looking at his phone, uttering a quantity expression in the form of a round number (e.g., "It's $200."). After seeing the dialogue between the characters, participants were asked to

determine which phone the speaker was looking at by selecting one of two images of a phone. In one image, the phone was turned face down, making the content of the screen invisible (**COVERED** screen). In the other image, the phone was turned face up with the display fully visible (**VISIBLE** screen), displaying a relevant number.

Two factors were manipulated: first, the persona of the displayed characters was varied (Nerdy vs. Chill). The presentation of the two types of characters was identical to the cartoons in Figure 1-2, with the question and the response provided as text in the speech bubble. Secondly, the relation of the number displayed on the visible phone to the one in the utterance was varied across two control and distractor levels and a critical one: *Match*, with full identity of 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, where the display could plausibly be seen as being close enough to have prompted the utterance, depending on the standard of precision adopted by the respondent. The Screen Fit manipulation is illustrated in Figure 6.
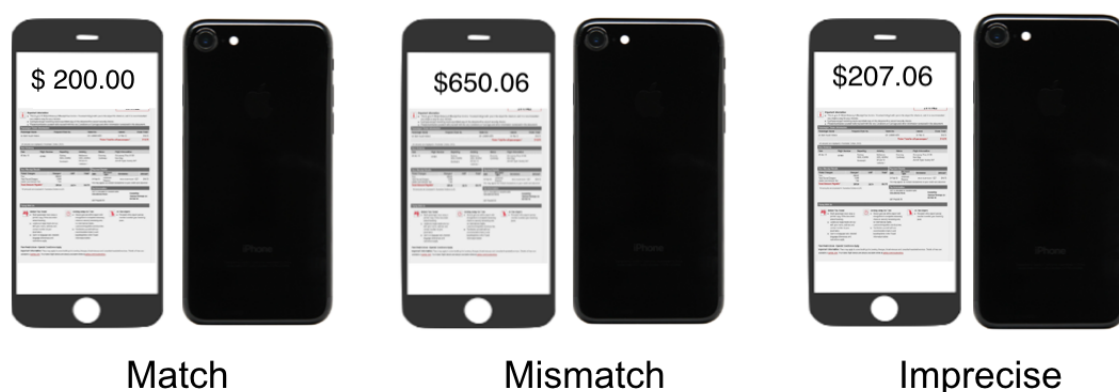
Figure 6: Screen Fit Manipulation

Participants were instructed to select the visible screen if they thought 'that the information on the screen fits what is being said;[6] and to select the covered screen if they believed it didn't. The Match and the Mismatch conditions serve as controls and are expected to evoke COVERED responses at floor and ceiling levels respectively, without any significant variation by speaker persona. By contrast, responses in the Imprecise condition crucially depend on the (im)precision standards employed by participants on a given trial. In particular, a strict interpretation – that is, one with a standard of precision that excludes the value displayed on the visible screen from the set of possible referents of the predicate – should lead to a COVERED screen response; and a more approximate interpretation – that is, one with a lower level of precision, which *includes* the value displayed on the visible screen – should translate into a VISIBLE screen choice. As outlined in §3 above, Hypothesis 1 predicts that the rate of covered screen choices will be higher for Nerdy speakers than Chill ones.

---

[6]See supplementary materials for full instructions.

Figure 7: Display before making the choice (Condition: Nerdy, Imprecise)

5.2 Materials

24 experimental items were created, each varied across 6 different conditions resulting from the $2{\times}3$ manipulation of the factors described above. The Persona manipulation was administered between-subjects to avoid raising participants' awareness of this manipulation, i.e., they would see either only dialogues between the Nerdy characters or between the Chill characters. The Screen Fit manipulation was administered within-subjects: each participant saw 6 items in both 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 1-5); 8 items contained utterances describing distances, expressed in miles; and 8 items contained items describing times, expressed in hours and minutes.[7] The experiment also included 24 filler items, which featured a dialogue between two separate characters. The fillers were alternated with the experimental items, so that participants would never see two consecutive fillers or experimental items.[8]

5.3 Procedure

The study was implemented and administered online on the PCIbex platform (see Schwarz and Zehr 2021).[9] For each item, the context sentence was introduced first on the top left of the screen. Sequentially, images of the two characters would appear, followed by the question asked by the first character and the answer from the other character. Next, the question task-prompt for participants

---

[7] Across the three types of numerals, the range of deviation in the Imprecise condition was always comprised between 5% and 18%. See supplementary materials for a full list of items.

[8] See supplementary materials for a list of filler items.

[9] https://www.pcibex.net

as well as the two pictures 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 respective phone pictures 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 or an obvious mismatch. Feedback was provided on the practice items, so as to help the participants familiarize themselves with the task. Figure 7 provides a full illustration of the display that participants would see once all the elements appeared.

5.4 Participants

189 self-declared native speakers of English (Age Median=31; female=129; male=56; non-binary=4) were recruited on Prolific and compensated $2 ($8/hour).

5.5 Results

Our first step is to ascertain the validity of our two control conditions – Match and Mismatch – which were set up to lead to ceiling and floor levels of COVERED choice rates. The proportions of COVERED choice rates across conditions is shown in Figure 8. The descriptive pattern suggests a clear step-wise effect of Screen Fit Manipulation, as expected, with ceiling and floor-level response rates for the controls, and the imprecise condition in the middle. To assess the validity of the two controls, we compared the rate of COVERED responses in the Imprecise condition with those in Mismatch and Match, collapsing across Persona levels in a mixed-effect logistic regression using the lme4 package in R. As expected, the rate of COVERED choices in the Imprecise condition was significantly higher than the rate in the Match (β=-7.8; SE=0.41; $z = 18.85$, p <0.001) but lower than the rate in the Mismatch condition (β=5.44; SE=0.29; $z = 19.04$, p <0.001).
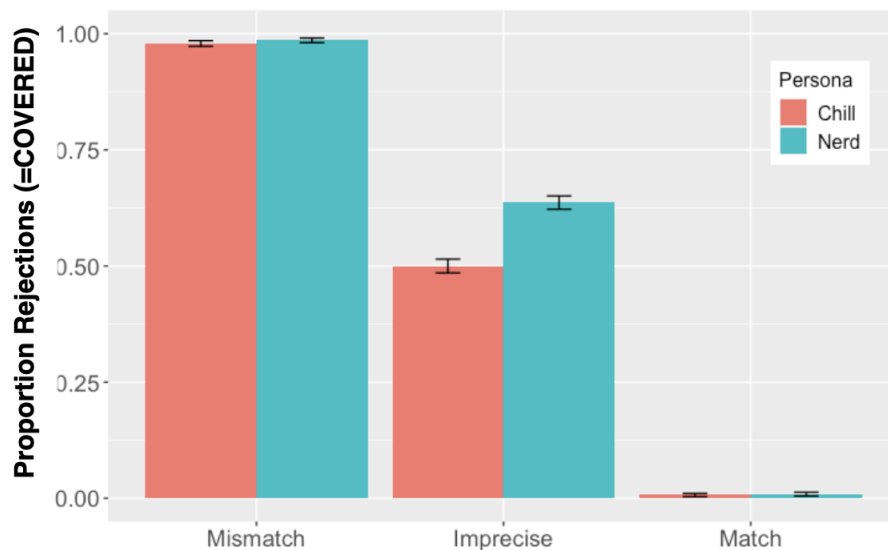

Figure 8: COVERED choices across Screen Fit, split by Persona

Next, we explore the effect of Persona in the critical Imprecise condition: as can be seen in Figure 8, we observe higher covered screen choice rates for Nerdy speakers compared to Chill

ones. We fit another mixed-effect logistic regression to the Imprecise condition data, with Persona as a fixed effect, and by-Subject and by-Item random intercepts.[10] This confirmed the significance of the Persona effect, with a higher rate of COVERED responses for Nerdy than for Chill speakers ($\beta = 1.54$; $SE = 0.52$; $z = 2.95$, $p < 0.01$).

5.6 Discussion

The results from Experiment 1 lend support to **Hypothesis 1**: numerals uttered by Nerdy speakers are associated with a narrower range of values than numerals uttered by Chill speakers, suggesting that the speaker persona impacted the computation of the extension of the numeral. At the same time, these findings cannot speak to how comprehenders engage in the reverse reasoning – i.e., how they go about assessing the appropriateness of a numerical description in light of a known, established value. As can be recalled from the discussion in §3, this aspect is another dimension central to the space of indeterminacy created by imprecision – and one that could be affected by social information in two different ways. In order to evaluate these possibilities, we now turn to Experiment 2, which uses a version of a Truth-Value judgment task.

6 Experiment 2: Assessing statements about given facts

In §3 we outlined two alternative hypotheses about how our Persona manipulation might affect comprehenders' adjudication of imprecise statements vis-a-vis an established fact. If comprehenders base their assessment on the range of values in the extension of the numeral – i.e., a narrower one for Nerdy speakers – imprecise descriptions coming from Nerdy speakers should be rejected as inappropriate more often than those coming from Chill speakers (**Hypothesis 2A**), paralleling the results pattern from Experiment 1. But if, in this type of task, comprehenders reason about persona-based information on an epistemic level, and thus utilize social information to calibrate the level of credibility that the speaker has, Nerdy speakers could be seen as more credible than Chill ones, leading comprehenders to *accept* their imprecise descriptions more often than for Chill speakers (**Hypothesis 2B**). We test these hypotheses in Experiment 2.

6.1 Methods, Design, Materials & Procedure

Experiment 2 adapts the setup of Experiment 1 to a paradigm in which participants are provided with the actual state of affairs in question, and are asked to assess whether a given statement correctly describes it. This type of setup corresponds to the *Truth Value Judgment* task, widely utilized in experimental semantics and pragmatics (Crain and McKee 1985; Noveck 2001; Papafragou and Musolino 2003; Doran et al. 2012 i.a.; see also Waldon and Degen 2020; Scontras and Pearl 2021 for recent critical re-visitations of the assumptions behind this methodology).

We implemented this setup by leaving the logic of Experiment 1 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

---

[10]The model constituted the maximally complex random effect structure that would converge. Given the floor and ceiling levels of response patterns in the Match and Mismatch conditions, and the clear absence of a Persona effect there, we home in directly on the Imprecise condition here, due to issues with modeling of interaction data at the extremes, with response proportions extremely near to 0 and 1 in the controls.

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. The full display of an item is reproduced in Figure 9.

Figure 9: Experiment 2 . Display before making the choice (Condition: Nerdy, Imprecise)

The Screen Fit manipulation, implemented within-subjects, varied the fit between the number on the phone and the uttered number, again with 3 levels (Match, Mismatch, Imprecise). The Persona manipulation, implemented between-subjects, varied the speaker persona, with Nerdy and Chill as levels. All experimental and filler items were the same as those deployed in Experiment 1.

6.2 Participants

190 participants were recruited on Prolific and compensated $2 ($8/hour) for participating (Age range: 18-66; Age Mean: 24; female=132; male=53; other=5).

6.3 Results

As can be seen in Figure 10, rejections (="wrong" responses) were at ceiling and at floor in the Mismatch and Match condition respectively, with intermediate 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).

As in Experiment 1, our main interest was to assess the effect of Persona in the critical, Imprecise condition. To explore this effect, we fit a mixed-effects logistic regression with Persona as fixed effect, and by-Subject and by-Item random intercepts.[11] The rejection rates for the two Per-

---

[11]Again, this model constituted the maximally complex random effect structure that would converge.

sona levels differed significantly from each other, but in contrast with Experiment 1, rejection rates were higher for Chill speakers than for Nerdy speakers (β=1.13; SE=0.52; $z = 2.18$, $p <$0.05).
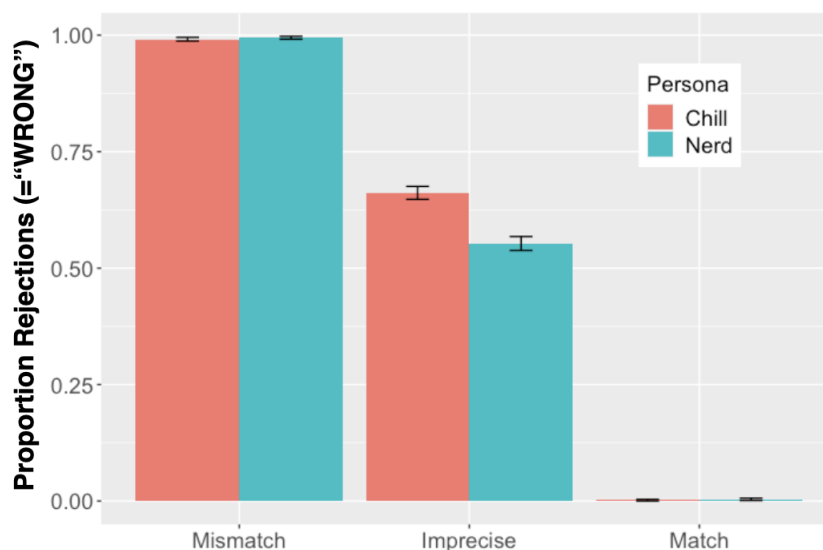

Figure 10: COVERED choices across Screen Fit, split by Persona

6.4 Combined analysis

In light of the opposite results in Experiments 1 and 2, and the fact that they only differed minimally in the specifics of the task, we pooled the data from the two experiments together and carried out a combined analysis. For a direct visual comparison, Figure 11 illustrates the rejection rates in the Imprecise condition across the two tasks – with rejection corresponding to a COVERED SCREEN choice in Experiment 1, and to a WRONG response in Experiment 2.
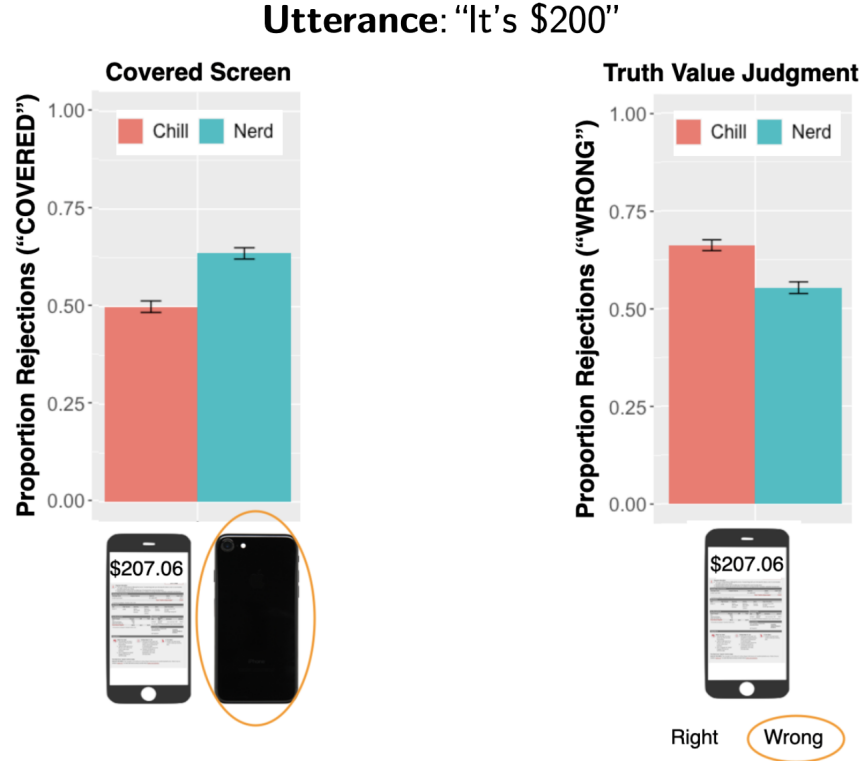
Figure 11: Rejection rates across tasks: Covered Screen (Exp1) vs. Truth Value Judgement (Exp2)

To compare these rates, we fit a mixed effect logistic regression model with Task and Persona as interacting and centered fixed effect factors and by-Item and by-Subject random intercepts (the most complex random effect structure resulting in a converging model). The model revealed a significant interaction (β=2.70; SE=0.73; $z = 3.69$, p <0.001), and no significant main effects (Persona: β=0.38; SE=0.34; $z = 1.11$, $p = 0.26$; Task: β=0.55; SE=0.34; $z = 1.60$, $p = 0.10$). This provides direct statistical confirmation that the impact of Persona has different, and indeed opposite, effects in the two task variants.

6.5 Discussion

The findings from Experiment 2 support **Hypothesis 2B**: rather than basing their adjudication of imprecise numerals on the range of values expected to fall under the numeral's extension given the speaker's identity, comprehenders appear to reason about persona-based information to calibrate the degree of epistemic credibility that they are willing to ascribe to the speaker, resulting in a more charitable behavior towards Nerdy speakers than Chill ones – and correspondingly, a rejection pattern opposite from the one found in Experiment 1.

Before considering the broader implications of this result, a lingering question needs to be addressed: *why* do we observe this pattern, as opposed to the alternative one outlined in Hypothesis 2A (see §3) – i.e., that Nerdy speakers' descriptions, by virtue of being associated with a more narrow range of values, should actually be rejected more often? We suggest that the relevant distinctive property of the task deployed in Experiment 2 is that in assessing whether a statement is right or wrong, a rejection response is crucially *prejudicial*: it commits the comprehender to

holding the speaker *blameworthy* for not abiding to the norms of conversation in terms of using language untruthfully. This contrasts with the picture selection paradigm in Experiment 1, where a participant's selection of the COVERED screen simply indicates that the character in the picture must have been looking at a different phone – making a rejection response entirely compatible with the assumption that the speaker is behaving in a cooperative fashion. We suggest that it is precisely the prejudicial nature of a rejection in Experiment 2 that tipped the scales against the response pattern predicted by Hypothesis 2A, 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 conversational norms than Chill ones. In other words, the expectation of higher precision for Nerdy speakers goes hand in hand with a higher level of credibility assigned to them, generating a more charitable disposition towards these speakers when it comes to assessing the appropriateness of their statements to represent a fact.

7   General Discussion

The key finding from our study is that comprehenders engaging in pragmatic reasoning utilize information about the social persona embodied by the speaker in two different ways. In Experiment 1, numerals were interpreted more precisely – i.e., as describing a narrower range of values – when uttered by Nerdy characters, a type of speaker who is expected to describe things precisely – as opposed to Chill ones – a type of speaker who is expected to describe things less precisely. In contrast, Experiment 2 showed that imprecise numeral descriptions by Nerdy characters are accepted as 'right' more often than those by Chill speakers. We argue that this shows that comprehenders recruit persona-based information to reason not only on the descriptive level, but also on the epistemic one – that is, to determine how much credibility they are willing to grant to someone when it comes to assessing whether they are misrepresenting the relevant facts. We now turn to discussing the implications of these findings, focusing on three major points.

First, our findings highlight comprehenders' social perception of speakers as a factor shaping the overall pragmatic inferencing that comprehenders engage when interpreting the content of an utterance. This contributes to a more inclusive view of pragmatic reasoning as a socially embedded process, adding to recent work demonstrating the impact of social considerations broadly construed on meaning-related inferences, e.g. relating to politeness (Bonnefon and Villejoubert 2006; Bonnefon et al. 2009; Mazzarella et al. 2018; Zhang and Wu 2020; Yoon et al. 2016, 2020), native speaker status (Fairchild and Papafragou 2018, and political orientation (Mahler 2020, 2022; Henderson and McCready 2019; see discussion in §1). In particular, our results shed a new light on the interaction between social information and meaning interpretation: they highlight the role of persona-constructs for shaping this process, and thus point to a bi-directional link between the constellation of identity and personality features distinctive of a speaker and the resolution of imprecision – and meaning more broadly. Overall, this broadens our understanding of the types of contextual cues at play in interpretation, suggesting new parallels between pragmatic reasoning and other domains of linguistic cognition, such as phonetic processing (Strand 1999; Niedzielski 1999; Hay et al. 2006; Hay 2009; Babel 2012; Drager 2015; Sumner et al. 2014; D'Onofrio 2016, 2018; Wade 2022 i.a.; see §1 for details), similarly linked to how listeners track different sources of social information in the utterance context.

But our findings also carry implications for semantic and pragmatic theory more broadly, as

they raise the question of whether, and how, current models of pragmatic inference may need to be enriched to incorporate the role of persona stereotypes into processes of utterance interpretation. While providing an answer to this issue goes beyond the scope of our paper, we see recently developed *Rational Speech Act* models as a promising starting point to undertake this endeavor. As suggested in §1, this framework was designed to capture the reasoning process behind meaning resolution in probabilistic terms, drawing on Grice's view of communication as a goal-oriented, cooperative activity (Frank and Goodman 2012; Kao et al. 2014; Goodman and Frank 2016; see §1), and placing an emphasis on the idea that interlocutors engaging in conversation are aiming at maximizing their *utility* vis-a-vis the goals that they are pursuing (e.g., informing each other about the state of the world). Besides embracing a view of utterance interpretation as a branch of social cognition, fully consistent with our findings, this framework is especially appealing in that it has already been productively deployed to capture the two phenomena of central interest: (im)precision (Kao et al. 2014; Aparicio 2017; Aparicio et al. 2022), together with closely related phenomena such as vagueness and gradability (see Lassiter and Goodman 2013, 2017; Goodman and Frank 2016; Qing 2021; Xiang et al. 2022); and the construction and negotiation of social identity, mostly in connection with speech sounds (Burnett 2017, 2019). While these works have not addressed the connection between speaker identity and utterance interpretation explicitly, they make available two key insights as potential starting points. One is the idea that a speaker's choice to describe things more or less precisely stems from striking a balance between accuracy and effort – an idea already suggested in formal and experimental approaches to the study of (im)precision (Van Der Henst et al. 2002; Krifka 2007; Kao et al. 2014; Solt et al. 2017; Klecha 2018; Aparicio 2017 i.a.). The other key ingredient is the idea that socio-indexical information, including the speaker's identity and personality features, is part and parcel of the message conveyed by an utterance; and that, accordingly, interlocutors are not just pursuing the goal of describing reality, but also of projecting a certain persona successfully in interaction – which in turn affects their strategies to maximize utility (Burnett 2017, 2019). Accordingly, comprehenders might reason that speakers weigh accuracy versus effort in different ways, depending on what persona they wish to make salient: they will posit that speakers willing to project a Nerdy persona will put a premium on accuracy at the cost of incurring a higher effort; and that speakers willing to project a Chill persona will put a premium on effortlessness, at the cost of sacrificing accuracy. This reasoning, in turn, will lead them to adjust their interpretive behavior accordingly – both when it comes to determining the referent of the numeral, or to calibrating the charity towards a potentially misbehaving speaker. We suggest that this idea might be implemented by enriching RSA models to include different expectations about the accuracy/complexity tradeoff for speakers embodying different personae, and allowing these expectations to impact comprehenders' reasoning accordingly.

Another broader 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 in common meaning-based tasks, namely picture selection and Truth-Value judgments. In recent years, a growing body of work has called for a critical re-assessment of the nature and significance of such judgments, highlighting the frequent impact of various factors – and thus sources of variability – that go beyond the purely linguistic properties of a given utterance. These include, among others, the discourse context (Sikos et al. 2019; Scontras and Pearl 2021) and the probability with which respondents see themselves as producing the utterance (Waldon and Degen 2020). Our results further expand on the context-sensitivity of these types of behavioral measures, suggesting that comprehenders track an array of contextual factors going beyond those traditionally

investigated in pragmatics, including highly specific social constructs such as the speaker persona. Moreover, and rather intriguingly, the way that speaker identity impacts how imprecision is dealt with varies across the two tasks we looked at. While in the COVERED SCREEN task, the extent of the range that a given numeral is taken to describe aligns with expectations about a speaker's level of precision, the perceived correctness of an imprecise description in the truth-value judgment task is *inversely* correlated with the expectation that a speaker uses this description imprecisely. As we argue, this means 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 right or wrong. This complex interplay between social meaning and interpretive tasks calls for a more extensive consideration of how behavioral 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.

A third implication of our results is that the effect of persona-based information we observed on the interpretation of numerals presents parallels with processes outside the domain of linguistics proper – and in particular with how interacting agents determine the degree of authority and knowledge that they are willing to ascribe to someone on a given matter. As discussed in §1, it has been argued that this reasoning often results in prejudicial practices such as *credibility deficit* (Fricker 2007) and *credibility excess* (Davis 2016). Framed in the broader picture of how social stereotypes shape people's behavior, our results thus suggest that processes of utterance interpretation are likely to be subject to, and simultaneously reproduce, similar patterns of bias and prejudice: the higher charity granted to Nerdy speakers suggests that comprehenders treat them as more credible, and hence more trustworthy, than their Chill counterparts, and hence more deserving of being given the benefit of the doubt in a situation in which they are at risk of being deemed conversationally culpable. This points to yet another uncharted dimension of investigation for the study of utterance interpretation: similar mechanisms linking social identity, utterance interpretation and the ascription of credibility are likely to be at play – and likely more impactful – when more socially prejudicial dimensions of identity are involved – e.g., race, gender socioeconomic status – in line with what has been suggested in work on linguistic discrimination more broadly (see Lippi-Green 2011; Rosa 2016; Rickford and King 2016; Jones et al. 2019 among many others). This, in turn, highlights the importance of engaging in further work aimed at unpacking the link between meaning comprehension and the perception and evaluation of social identity to illuminate how these domains of reasoning jointly underlie people's communicative behavior.

8   Conclusion

We presented experimental evidence that comprehenders' interpretation of numerals is affected by the social persona embodied by the speaker. We believe our results provide a novel perspective for considering the relationship between pragmatic reasoning and the socio-indexical information available in the context; and that, accordingly, there is much to gain from developing a comprehensive approach to the study of how interlocutors navigate different sources of information to represent and communicate about the state of the world.

Appendix

A full list of experimental and filler items, instructions in full, data files, and R code for data analysis and visualization are available here:
`https://osf.io/4efqr/?view_only=028c2e00f37f46778d0410fed25f5748`.

References

Acton, E. (2019). Pragmatics and the social life of the english definite article. *Language 95*(1), 37–65.

Acton, E. and C. Potts (2014). That straight talk. Sarah Palin and the sociolinguistcs of demonstratives. *Journal of Sociolinguistics 18(1)*, 3–31.

Acton, E. K. (2022, jan). Sociophonetics, semantics, and intention. *Journal of Linguistics 58*(3), 465–494.

Agha, A. (2003). The social life of a cultural value. *Language and Communication 23*, 231–273.

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

Aparicio, H. (2017). *Processing context-sensitive expressions: the case of Gradable Adjectives and Numerals*. Ph. D. thesis, University of Chicago.

Aparicio, H., C. Chen, R. Levy, and E. Coppock (2022). Granularity in the semantics of comparison. *Semantics and Linguistic Theory 31*, 550.

Aparicio, H., M. Xiang, and C. Kennedy (2016). Processing gradable adjectives in context: A visual world study. *Semantics and Linguistic Theory 25*, 413–432.

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

Babel, M. (2012). Evidence for phonetic and social selectivity in spontaneous phonetic imitation. *Journal of Phonetics 40*(1), 177–189.

Beltrama, A. (2018). Precision and speaker qualities. the social meaning of pragmatic detail. *Linguistics Vanguard 4*(1).

Beltrama, A. (2020). Social meaning in semantics and pragmatics. *Language and Linguistics Compass 14*(9).

Beltrama, A., S. Solt, and H. Burnett (2022). Context, precision, and social perception: A socio-pragmatic study. *Language in Society*, 1–31.

Beltrama, A. and L. Staum Casasanto (2017). Totally tall sounds totally younger. intensifiers at the socio-semantic interface. *Journal of Sociolinguistics 21*(2), 154–182.

Beltrama, A. and L. Staum Casasanto (2022). The social meaning of semantic properties. In L. Hall-Lew, E. Moore, and R. J. Podesva (Eds.), *Social Meaning and Linguistic Variation: Theorizing the Third Wave*. Cambridge University Press.

Benor, S. B. (2004). 'talmid chachams and tsedeykeses': Language, learnedness, and masculinity among orthodox jews. *Jewish Social Studies 11*(1), 147–170.

Bergen, L. (2016). *Joint inference in pragmatic reasoning*. Ph. D. thesis, Massachusetts Institute of Technology.

Bonnefon, J.-F., A. Feeney, and G. Villejoubert (2009). When some is actually all: Scalar inferences in face-threatening contexts. *Cognition 112*(2), 249 – 258.

Bonnefon, J.-F. and G. Villejoubert (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.

Bott, L., T. M. Bailey, and D. Grodner (2012). Distinguishing speed from accuracy in scalar implicatures. *Journal of Memory and Language 66*(1), 123 – 142.

Bucholtz, M. (1999). "Why be normal?": Language and identity practices in a community of nerd girls. *Language in Society 28*(2), 203–223.

Bucholtz, M. (2001). The whiteness of nerds: Superstandard English and racial markedness. *Journal of Linguistic Anthropology 11*, 84–100.

Burnett, H. (2017). Sociolinguistic interaction and identity construction: The view from game-theoretic pragmatics. *Journal of Sociolinguistics 21*(2), 238–271.

Burnett, H. (2019). Signalling games, sociolinguistic variation and the construction of style. *Linguistics and Philosophy 42*, 419–450.

Campbell-Kibler, K. (2010). Perception in sociolinguistics. *Language and Linguistics Compass. 4(6)*, 377–389.

Campbell-Kibler, K. (2011). The sociolinguistic variant as a carrier of social meaning. *Language Variation and Change 22*, 423–441.

Choe, J., S. Sloggett, M. Yoshida, and A. D'Onofrio (2019). Personae in syntactic processing: Socially-specified agents bias expectations of verb transitivity. Poster presented at CUNY 2019.

Coupland, N. (2007). *Style: Language Variation and Identity*. Cambridge University Press.

Crain, S. and C. McKee (1985). Acquisition of structural restrictions on anaphora. *North East Linguistics Society (NELS) 16*(1).

Cummins, C., U. Sauerland, and S. Solt (2012). Granularity and scalar implicature in numerical expressions. *Linguistics and Philosophy 35*, 135–169.

Davis, E. (2016). Typecasts, tokens, and spokespersons: A case for credibility excess as testimonial injustice. *Hypatia 31*(3), 485–501.

de Jong, K., M. E. Beckman, and J. Edwards (1993). The interplay between prosodic structure and coarticulation. *Language and Speech 36*(2-3), 197–212.

Degen, J. and M. K. Tanenhaus (2015). Processing scalar implicature: a constraint-based approach. *Cognitive science 39*(4), 667–710.

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

D'Onofrio, A. (2016). *Social meaning in linguistic perception*. Ph. D. thesis, Stanford University.

D'Onofrio, A. (2018). Personae and phonetic detail in sociolinguistic signs. *Language in Society 47*(4), 513–539.

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

D'Onofrio, A. (2020). Personae in sociolinguistic variation. *WIREs Cognitive Science 11*(6).

Doran, R., G. Ward, M. Larson, Y. McNabb, and R. E. Baker (2012). A novel experimental paradigm for distinguishing between what is said and what is implicated. *Language 88*(1), 124–154.

Drager, K. K. (2015). *Linguistic variation, identity construction and cognition*. Language Science Press.

Eckert, P. (2000). *Language variation as social practice*. Oxford: Blackwell.

Eckert, P. (2008). Variation and the indexical field. *Journal of Sociolinguistics. 12:4*, 453–76.

Fairchild, S. and A. Papafragou (2018). Sins of omission are more likely to be forgiven in non-native speakers. *Cognition 181*, 80–92.

Fiske, S. T. (2018). Stereotype content: Warmth and competence endure. *Current Directions in Psychological Science 27*(2), 67–73.

Frank, M. C. and N. D. Goodman (2012). Predicting pragmatic reasoning in language games. *Science 336*(6084), 998–998.

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

Gal, S. and J. T. Irvine (2019). *Signs of Difference: Language and Ideology in Social Life*. Cambridge University Press.

Gazdar, G. (1980). Pragmatics and logical form. *Journal of Pragmatics 4*(1), 1–13.

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

Goodman, N. D. and M. C. Frank (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences 20*(11), 818 – 829.

Goodman, N. D. and A. Stuhlmüller (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science 5*(1), 173–184.

Grice, H. P. (1975). Logic and conversation. In P. Cole and J. L. Morgan (Eds.), *Syntax and Semantics, Vol. 3, Speech Acts*, pp. 41–58. New York: Academic Press.

Hall-Lew, L., E. Moore, and R. J. Podesva (Eds.) (2021). *Social Meaning and Linguistic Variation*. Cambridge University Press.

Hay, J. (2009). Stuffed toys and speech perception. *Linguistics 11*(3), 135–163.

Hay, J., P. Warren, and K. Drager (2006). Factors influencing speech perception in the context of a merger-in-process. *Journal of Phonetics 34*, 458–484.

Henderson, R. (2019, jun). Dogwhistles and the at-issue/non-at-issue distinction. In *Secondary Content*, pp. 222–245. BRILL.

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

Horn, L. (1984). Towards a new taxonomy of pragmatic inference: Q-based and r-based implicature. In *Meaning, form, and use in context: Linguistic applications*, pp. 11–42. Washington: Georgetown University Press.

Huang, Y., E. Spelke, and J. Snedeker (2013). What exactly do number words mean? *Language Learning and Development 9*(2), 105–129.

Huang, Y. T. and J. Snedeker (2018). Some inferences still take time: Prosody, predictability, and the speed of scalar implicatures. *Cognitive Psychology 102*, 105 – 126.

Hunt, M. and E. K. Acton (2022). "How's the wife?": Pragmatic reasoning in spousal reference. *Journal of Pragmatics 188*, 152–170.

Irvine, J. . (2001). Style as distinctiveness: the culture and ideology of linguistic differentiation. In J. Rickford and P. Eckert (Eds.), *Stylistic Variation*. Cambridge University Press.

Jenkins, A. C., P. Karashchuk, L. Zhu, and M. Hsu (2018). Predicting human behavior toward members of different social groups. *Proceedings of the National Academy of Sciences 115*(39), 9696–9701.

Jeong, S. (2021). Deriving politeness from an extended lewisian model: The case of rising declaratives. *Journal of Pragmatics 177*, 183–207.

Jones, T., J. Kalbfeld, R. Hancock, and R. Clark (2019, 01). Testifying while black: An experimental study of court reporter accuracy in transcription of african american english. *Language*.

Kao, J. T., J. Y. Wu, L. Bergen, and N. D. Goodman (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences 111*(33), 12002–12007.

Karawani, H. and B. Waldon (2017). Strength, duality and social meaning. *Amsterdam Colloquium 21*.

Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy 30(1)*(1), 1–45.

Kiesling, S. F. (2004, aug). DUDE. *American Speech 79*(3), 281–305.

Kiesling, S. F. (2018). Masculine stances and the linguistics of affect: on masculine ease. *NORMA 13*(3-4), 191–212.

Klecha, P. (2018). On unidirectionality in precisification. *Linguistics and Philosophy 41*(1), 87–124.

Krifka, M. (2007). Approximate interpretations of number words: A case for strategic communication. In G. Bouma, I. Krämer, and J. Zwarts (Eds.), *Cognitive foundations of interpretation*, pp. 111–126. Amsterdam: Royal Netherlands Academy of Arts and Sciences.

Labov, W. (1966). *The Social Stratification of English in New York City*. Washington, D.C.: Center for Applied Linguistics.

Lasersohn, P. (1999). Pragmatic halos. *Language 75(3)*(3), 522–551.

Lassiter, D. and N. D. Goodman (2013). Context, scale structure, and statistics in the interpretation of positive-form adjectives. In *Proceedings of SALT 23*.

Lassiter, D. and N. D. Goodman (2017). Adjectival vagueness in a bayesian model of interpretation. *Synthese 194*(10), 3801–3836.

Leffel, T., M. Xiang, and C. Kennedy (2016). Imprecision is pragmatic: Evidence from referential processing. *Semantics and Linguistic Theory (SALT) 26*.

Levinson, S. C. (2000). *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, MA: MIT Press.

Levon, E. (2016). Gender, interaction and intonational variation: The discourse functions of high rising terminals in london. *Journal of Sociolinguistics 20*(2), 133–163.

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

Lindblom, B. (1990). Explaining phonetic variation: A sketch of the h&h theory. In *Speech Production and Speech Modelling*, pp. 403–439. Springer Netherlands.

Lippi-Green, R. (2011). *English with an Accent*. Taylor & Francis Ltd.

Mahler, T. (2020). The social component of the projection behavior of clausal complement contents. *Linguistic Society of America 5*(1), 777.

Mahler, T. (2022). *Social identity information in projection inferences: a case study in social and semantic-pragmatic meaning*. Ph. D. thesis, The Ohio State University.

Mason, M., A. Lee, E. Wiley, and D. Ames (2013). Precise offers are potent anchors: Conciliatory counteroffers and attributions of knowledge in negotiations. *Journal of Experimental Social Psychology 49*, 759–763.

Mazzarella, D., E. Trouche, H. Mercier, and I. Noveck (2018). Believing what you're told: Politeness and scalar inferences. *Frontiers in Psychology 9*, 908.

Milroy, L. and D. R. Preston (1999). Introduction. *Journal of Language and Social Psychology 18*(1), 4–9.

Niedzielski, N. (1999). The effect of social information on the perception of sociolinguistic variables. *Journal of Social Psychology (Special Edition) 18:1*, 62–85.

Noveck, I. (2001). When children are more logical than adults: experimental investigations of scalar implicature. *Cognition 78*(2), 165 – 188.

Ochs, E. (1992). Indexing gender. In A. Duranti and C. Goodwin (Eds.), *Rethinking Context: Language as an Interactive Phenomenon*, pp. 335–358. New York, Cambridge: Cambridge University Press.

Papafragou, A. and J. Musolino (2003). Scalar implicatures: experiments at the semantics–pragmatics interface. *Cognition 86*(3), 253 – 282.

Pena-Marin, J. and R. Bhargave (2016). Lasting performance: Round numbers activate associations of stability and increase perceived length of product benefits. *Journal of Consumer Psychology 26*(3), 410 – 416.

Pinkal, M. (1995). *Logic and lexicon*. Dordrecht.

Podesva, R. J. (2007). Phonation type as a stylistic variable: The use of falsetto in constructing a persona. *Journal of Sociolinguistics 11.4*, 478–504.

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

Podesva, R. J., J. Reynolds, P. Callier, and J. Baptiste (2015). Constraints on the social meaning of released /t/: A production and perception study of u.s. politicians. *Language Variation and Change 27*, 59–87.

Qing, C. (2021). Zero or minimum degree? rethinking minimum gradable adjectives. *Proceedings of Sinn und Bedeutung*, Vol 25 (2021): Proceedings of Sinn und Bedeutung 25.

Rickford, J. R. and S. King (2016). Language and linguistics on trial: Hearing rachel jeantel (and other vernacular speakers) in the courtroom and beyond. *Language 92*(4), 948–988.

Roberts, C. (1996/2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics 5*(6), 1–69.

Rosa, J. (2016, nov). From mock spanish to inverted spanglish. In *Raciolinguistics*, pp. 65–80. Oxford University Press.

Sauerland, U. and P. Stateva (2007). Scalar vs. epistemic vagueness: Evidence from approximators. In M. Gibson and T. Friedman (Eds.), *Proceedings of Semantics and Linguistic Theory (SALT) 17*, Ithaca, NY. CLC Publications.

Sauerland, U. and P. Stateva (2011). Two types of vagueness. In *Vagueness and Language Use*, pp. 121–145. London: Palgrave Macmillan.

Schwarz, F., J. Romoli, and C. BIll (2016). Reluctant Acceptance of the Literal Truth: Eye Tracking in the Covered Box Paradigm. In N. Bade, P. Berezovskaya, and A. Schoeller (Eds.), *Sinn und Bedeutung (SuB) 20*, Tübingen, pp. 61–78. Uiversity of Tuebingen.

Schwarz, F. and J. Zehr (2021). Tutorial: Introduction to PCIbex – An Open-Science Platform for Online Experiments: Design, Data-Collection and Code-Sharing. *Annual Meeting of the Cognitive Science Society (CogSci) 43*(43).

Scontras, G. and L. S. Pearl (2021, oct). When pragmatics matters more for truth-value judgments: An investigation of quantifier scope ambiguity. *Glossa: a journal of general linguistics 6*(1).

Sikos, L., M. Kim, and D. J. Grodner (2019). Social context modulates tolerance for pragmatic violations in binary but not graded judgments. *Frontiers in Psychology 10*.

Silverstein, M. (2003). Indexical order and the dialectics of sociolinguistic life. *Language and communication 23*, 193–229.

Solt, S. (2014). An alternative theory of imprecision. *Semantics and Linguistic Theory (SALT) 24*.

Solt, S., C. Cummins, and M. Palmovic (2017). The preference for approximation. *International Review of Pragmatics 9*(2), 248 – 268.

Stolier, R., H. E., and F. J.B. (2020). Trait knowledge forms a common structure across social cognition. *Nat Hum Behav 4*, 361–371.

Strand, E. A. (1999). Uncovering the role of gender stereotypes in speech perception. *Journal of Language and Social Psychology 18*(1), 86–100.

Sumner, M., S. K. Kim, E. King, and K. B. McGowan (2014). The socially weighted encoding of spoken words: a dual-route approach to speech perception. *Frontiers in Psychology 4*.

Syrett, K. and A. 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.

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

Thomas, W. C. (2021). Social inferences from the use of just as an exclusive particle. *Linguistic Society of America 6*(1), 746.

Van Der Henst, J.-B., L. Carles, and D. Sperber. (2002). Truthfulness and relevance in telling the time. *Mind and Language 17*, 457–466.

Wade, L. (2022). Experimental evidence for expectation-driven linguistic convergence. *Language*.

Waldon, B. and J. Degen (2020). Modeling behavior in truth value judgment task experiments. *Society for Computation in Linguistics (SCL) 3*(3).

Welsh, M., D. Navarro, and S. Begg (2011). Number preference, precision and implicit confidence. In *Annual Meeting of the Cognitive Science Society (CogSci) 33*, pp. 1521–1526.

Xiang, M., W. Xu, C. Kennedy, and T. Leffel (2022). Pragmatic reasoning and semantic convention: A case study on gradable adjectives. To Appear in Semantics & Pragmatics.

Xie, G.-X. and A. Kronrod (2012). Is the devil in the details? *Journal of Advertising 41*(4), 103–117.

Yoon, E. J., M. H. Tessler, N. D. Goodman, and M. C. Frank (2016). Talking with tact : Polite language as a balance between kindness and informativity.

Yoon, E. J., M. H. Tessler, N. D. Goodman, and M. C. Frank (2020). Polite speech emerges from competing social goals. *Open Mind 4*, 71–87.

Zhang, J. and Y. Wu (2020). Only *youxie* think it is a nice thing to say: Interpreting scalar items in face-threatening contexts by native Chinese speakers. *Journal of Pragmatics 168*, 19–35.

Zhang, Y. C. and N. 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.