

(Im)precise personae: The effect of socio-indexical meaning on semantic interpretation

Abstract. In this paper, we address the following question: how do comprehenders reason about the *persona* embodied by the speaker to determine the referential meaning of numerical expressions such as “The price is \$200”? Using a picture selection task, we show that descriptions uttered by speakers embodying a nerdy persona, indexically associated with highly precise speech, are interpreted more precisely than those uttered by speakers embodying a Chill persona, indexically associated with imprecise speech; and that this effect is especially prominent for comprehenders who do not see themselves as embodying the distinctive social qualities of the speaker. These findings enrich sociolinguistic theory in two ways: they contribute to building a more integrative perspective on the interface between socio-indexical meaning and other domains of signification — particularly at the *referential level*; and they further our understanding of how socio-indexical meaning affects linguistic cognition beyond the phonetic and morpho-syntactic domain.

Keywords: socio-indexical meaning; pragmatic reasoning; persona; imprecision; numerals

1 Introduction

Linguistic utterances carry two distinct kinds of meaning: a *descriptive*, or referential, meaning – i.e., what an utterance says about the world; and a *socio-indexical* meaning – i.e., what an utterance conveys about speakers themselves – e.g., their demographic background, ideological orientation, and personality traits (Labov 1966; Ochs 1992; Silverstein 2003; Eckert 2008 i.a.). Scholars across linguistic subfields have long acknowledged the central role of both components in linguistic com-

munication (Ochs 1992; Duranti 1997; Eckert 2019), highlighting how each of them enable conversational participants to pursue two key activities in linguistic interaction: exchanging propositional information, and thus zeroing in on the way things are in the world; and projecting and perceiving interlocutors' identity and personality, thus reflecting and reshaping the social context in which interaction takes place. At the same time, much remains to be discovered about the interaction between these layers of meaning, and in particular, on how interlocutors triangulate between them to assemble, negotiate and ascribe legible content in linguistic communication. This open question, we suggest, can be partly traced back to the fact that these two domains of signification have been by-and-large investigated in separate scholarly traditions.

On the one hand, socio-indexical meaning has been primarily explored by scholars in sociolinguistics and linguistic anthropology, whose work has framed axes of social differentiation between speakers, and the way in which they come to be indexed by linguistic forms, as a major force underlying patterns of linguistic variation and interaction (Silverstein 2003; Eckert 2008; Campbell-Kibler 2011; Podesva 2011; Levon 2014 among many others). Among the insights provided by this research, two are especially relevant. One is that social meanings do not boil down to lists of individual demographic and personality traits about the speaker, but rather come together in broader, holistic social constructs that index *types* of speakers, or *personae* – e.g., “Valley Girls”, “Business Professionals”. “Surfer Dudes” (Irvine 2001; Agha 2005; Coupland 2007; Eckert 2000; Podesva 2011; Kiesling 2018; D’Onofrio 2018; King 2021). It is sometimes at the level of these constructs that social meanings become maximally salient in the social space, and thus provide interlocutors with a powerful resource for projecting and perceiving identity and stylistic categories in interaction. Second, social meanings crucially interact with the psychological dynamics behind language processing, providing a cue for comprehenders to parse linguistic input across different domains – e.g., sound categorization and recognition; convergence; imitation. Effects of this sort have been demonstrated both in connection to demographic dimensions such as origin, age and gender (Strand 1999; Niedzielski 1999; Hay et al. 2006; Hay 2009; Babel 2012; Drager 2015; Staum Casasanto 2008; Sumner et al. 2014; Wade 2020); but also to specific personae. For

instance, listeners primed with social types such as “Valley Girl” displayed a backed perception of the vowel in TRAP and LOT than listeners primed with other social types (D’Onofrio 2018), suggesting a link in the cognitive representation of this particular persona and /a/ retraction (see also D’Onofrio 2015, 2016; Choe et al. 2019 for further work; and D’Onofrio 2020 for an overview).

On the other hand, work in semantics and pragmatics has treated socio-indexical information as generally only tangential to how descriptive meanings are computed. These approaches crucially posit that descriptive meanings are also, to a certain extent, indeterminate; as a result, they frame meaning interpretation not as a passive act of decoding, but as a dynamic process in which conversational participants play an active role. For instance, interlocutors can determine whether quantifiers such as *some* should be interpreted as having, or lacking, an upper-bounded interpretation (i.e., “some, but not all”) by reasoning on elements in the conversational setting – what question is being addressed in discourse, what the speaker’s communicative intention is, and what specific body of world knowledge can illuminate these assessments (e.g. Grice 1975; Horn 1984; Roberts 2012 i.a.). Yet, even in the face of this space of interpretive agency, considerations about conversational participants’ social identity have typically been seen as peripheral to how such inferences are drawn. This separation is especially clear in Grice’s (1975) foundational theory of communication, which has inspired much work of how descriptive meanings are ascribed: in this system, as well as in its later developments, conversational partners are effectively treated as socially undifferentiated interlocutors – i.e., idealized rational agents who, keeping general contextual factors constant, can be expected to resolve the meaning of an utterance in the same way, independent of how they might differ from each other in the social space.¹

In recent years, a growing body of research has spearheaded a more integrated approach to the study of meaning across the descriptive and socio-indexical dimension, highlighting a variety of domains in which the two interact. On the one hand, comprehenders have been shown to promptly infer social meanings from the semantic and pragmatic properties of their utterances, and across a variety of linguistic phenomena: demonstratives and determiners (Acton and Potts 2014; Acton 2019; Hunt and Acton 2022); intensifiers (Beltrama and Staum Casasanto 2017); modals

(Glass 2015); discourse particles (Thomas 2021); speech acts (Jeong 2021; see Beltrama 2020 for an overview). On the other hand, work on the resolution and comprehension of descriptive meanings has started to investigate the role of previously neglected factors in the process – e.g., politeness considerations (Bonneton et al. 2009; Yoon et al. 2020; Mazzearella et al. 2018; Zhang and Wu 2020); affect (Kao et al. 2014; Bergen 2016); and speaker-specific information such as linguistic nativeness (Fairchild and Papafragou 2018) or political orientation (Mahler 2020). Accordingly, theories modeling pragmatic inferences have been broadened accordingly, suggesting that the dynamics whereby interlocutors exchange propositional information can be conceptualized as effectively similar to those whereby they signal and infer social meanings conveyed by phonological variables (Burnett 2017, 2019).

But even in the face of these developments, much remains to be seen on the interplay between the wealth of nuanced, multi-dimensional socio-indexical associations that are available to interlocutors in interaction; and the reasoning dynamics whereby comprehenders resolve, and ultimately compute, the descriptive meaning of linguistic forms. In this respect, the demonstrated effects of social information on sound processing naturally raises the hypothesis that the very same clusters of speakers' characteristics and qualities – and the personae in which they coalesce – are likely to affect other domains of linguistic cognition – including that pertaining to how descriptive meanings are resolved and ascribed. We see the endeavor of exploring this hypothesis as important for two distinct, related reasons. First, it would contribute the opportunity of exploring whether, and how, socio-indexical information informs pragmatic reasoning and the ascription of referential meaning, contributing to a more comprehensive view of how different domains of signification enable interlocutors to exchange content in communication. Second, it would open up a novel testing ground to understand how indexical associations shape the cognitive processes behind language comprehension beyond the domain of phonological variation. We now turn to discuss the case study that we rely on to address this issue: *(im)precision* in the interpretation of numerals.

2 Imprecision, personae, interpretation: implementation and hypotheses

2.1 Imprecision: a socially meaningful space of variation

Consider the following example:

- (1) The ticket costs \$300.

Although prices and other quantity expressions seem to denote specific values, they are routinely used in a more liberal way: it is generally reasonable for someone to utter (1) when the actual price is \$295, suggesting that comprehenders can apply some deviation from the literal truth conditions when interpreting numerals. This phenomenon, known as *imprecision*, has been extensively investigated in work in semantics and pragmatics (Lewis 1979; Pinkal 1995; Lasersohn 1999; Solt 2014; Aparicio 2017 i.a.), emerging as an ideal testbed to investigate the interplay between social meanings and semantic interpretation for two reasons.

One property is that the possibility of speaking imprecisely introduces a space of indeterminacy around the interpretation of quantity expressions, such that comprehenders have to actively reason about the conversational setting to compute their descriptive meaning: “\$300”, for instance, can be taken to represent different price ranges – e.g., the exact price of \$300; the \$295-\$305 interval; or the larger \$290-\$310 interval. Only by tracking the specifics of the communicative setting can interlocutors settle on the extent of imprecision that should be applied, and hence zero in on what facts numerical expressions can be taken to describe (see Van Der Henst et al. 2002; Solt et al. 2017; Aparicio 2017 for work investigating this reasoning process).

The second property is that variation in precision is socially meaningful. In particular, speakers using sharp numbers (e.g., 203) – normally taken to signal a high level of precision – are perceived as embodying social qualities pertaining to high status and intellectual standing – e.g., being articulate, intelligent, educated, hard-working – as well as low solidarity and likability – e.g., annoying, pedantic, obsessive, and uptight; conversely, speakers using numerals in an explicitly imprecise fashion (i.e., “approximately 200”) have been shown to be perceived as embodying

opposite sets of qualities – e.g., likable, laid-back, friendly– and in a way that is remarkably robust across conversational contexts (Beltrama 2018; Beltrama et al. 2022). This suggests that different levels of precision are taken by comprehenders to index distinct constellations of socio-indexical information, similar to what happens with countless other examples of linguistic variation.

Taken together, these two properties of precision – its socio-indexical richness, and its semantic indeterminacy – provide an ideal testbed to illuminate how social meanings affect the ascription of descriptive meaning. More specifically: how do comprehenders reason about socio-indexical information to determine the level of precision for interpreting a numerical description, and thus zero in on its referent? We address this question experimentally, deploying a task that allows us to tap into comprehenders’ interpretation of numeral utterances produced by two distinct *types* of speakers: one embodying the social qualities typically perceived as indexed by high precision; and one embodying those linked to a lower degree of precision. To implement this contrast, we draw on the notion of speaker *persona* – a construct especially pertinent to our goals in two respects. First, as discussed in §1, personae are not just powerful semiotic resources for projecting and perceiving identity, they also shape language processing across different phenomena, thus emerging as a suitable dimension of social information to test the hypothesis that imprecision-related reasoning is similarly affected by identity-level distinctions between interlocutors. Second, persona-level constructs are often described with widely shared and recognizable labels and properties, and indexed by a wealth of both linguistic and non-linguistic signs (e.g., smoking, clothing etc.; see Eckert 2008 for further discussion), which makes them highly amenable to being made salient in an experimental setting. Accordingly, previous experimental work has successfully utilized these constructs, evoking them in different ways – e.g., a textual description of the persona at stake (D’Onofrio 2018); the display of objects stereotypically associated with the persona (e.g., a shopping bag for Valley Girls, see D’Onofrio 2015); or pictures of people embodying a particular persona (D’Onofrio 2019).

To address our question, we proceed in three steps. First, we implement (and norm) the persona contrast of interest (§2.2); second, we articulate our hypotheses concerning the effect of persona-

based information on numeral interpretation (§2.3); finally, we report on an experimental study that tests these hypotheses (§3).

2.2 The Persona contrast: implementation and norming

While pragmatic precision has not been explicitly linked to specific personae in previous work, the social meanings of this variable very much align with those invoked by similar variations along detail-orientedness in the phonetic domain. Specifically, work on the indexicality of *hyper-articulation*, a domain in which precision involves different phonetic components of an utterance, has unveiled a link between hyper-articulated speech and individual social qualities very similar to those evoked by precise numerals, including *articulateness* and *learnedness* (Bucholtz 2001; Benor 2004); *effortfulness* (Eckert 2008); and *detail-orientedness* (Podesva et al. 2015). Moreover, some of these social traits have been argued to coalesce in specific personae, foremost the *Nerds* in a California high school (Bucholtz 1999, 2001), associated with an emphasis on learnedness and educatedness, and on qualities overtly divorced from mainstream likability such as pedantry and un-coolness. In light of these considerations, we see a *Nerdy* persona as one that, following the broader associative pattern between detail-orientedness and speaker qualities discussed above, should be associated with the precise use of numerals as well.

As for the persona indexed by a lower degree of precision, we adopt a fundamentally *contrastive* approach: instead of searching for an elusive and perhaps non-existent instantiation of a pragmatically “neutral” baseline, we look to maximize the contrast between Nerdy speakers and speakers embodying a persona that can reasonably be expected to be associated with imprecision. To zero in on this persona, we draw on the observation that, similarly to the qualities associated with high precision, those linked to low precision/approximation – i.e., laid-backness, friendliness, coolness – have also been suggested by previous work to coalesce in specific, largely overlapping, persona constructs such as “Surfer dudes”, “Skaters”, and “Frat Boys” (Kiesling 2018) – a set of distinct types which nevertheless share a common indexical core of effortless, laid-backness and chillness, in line with the qualities indexed by imprecise speech.

To implement this contrast in the visual presentation of speakers, we created cartoon images of two characters engaged in conversation, indicated by the use of empty speech bubbles. One cartoon involved two characters, Arthur and Rachel, expected to embody the social qualities linked to high precision; the other involved two characters, Alex and Eva, expected to invoke those indexed by low precision. The two sets of characters are displayed in Figure 1.

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

To ascertain the viability of the persona contrast implemented above, we conducted a norming study comparing the social evaluation of the two sets of characters in Figure 1. 240 participants recruited on Prolific were shown these vignettes in a between-subject design: one half were shown Arthur and Rachel; the other half were shown Alex and Eva. Participants were asked to perform two tasks. In the first task, participants were asked to list three attributes and a stereotypical label for the characters shown. Participants' responses are illustrated in the word clouds in Figure ??-??, and show that the perception of Arthur and Alex largely aligns with what we aimed for: Arthur is overwhelmingly seen as embodying social qualities indicative of high intellectual standing (e.g., clever, smart) and is consistently associated with a Nerdy/Geeky persona. In contrast, Alex is ascribed attributes pertaining to chillness – e.g., laid-back, relaxed, easy, and cool; and high solidarity/sociability (e.g., friendly; outgoing). In the remainder of the paper, we refer to the two characters as the *Nerdy persona* and the *Chill persona* respectively, as a way of capturing the qualities that emerged as especially prominent for each of them.

Figure 2. Qualities ascribed to Nerdy vs. Chill characters in norming study.

In the second task, participants were asked to rate how precisely they thought the character they saw would speak about times and quantities on a 1-10 scale (10=maximally precise). As shown in Figure ??, the average precision ratings for Arthur were significantly higher than those for Alex ($t(238.67) = 5.63; p < 0.001$), in line with our expectation.

Figure 3. Expected precision for Chill vs/ Nerdy characters.

2.3 Personae and (im)precision resolution: hypotheses

We can now utilize this contrast to explore how comprehenders' perception of the speaker's persona affects how they interpret numerals. We are interested in two questions in particular. First, we ask whether, and how, the social persona embodied by the speaker affects the precision level required to interpret a numerical description, and therefore the computation of the range of values that this description can be taken to refer to. We hypothesize that an utterance produced by a speaker who embodies the qualities linked to high precision – a *Nerdy* one – will be associated with a higher standard of precision than the same utterance uttered by a speaker embodying the qualities linked to low precision – i.e., the *Chill* one. Accordingly, numerals uttered by Nerdy speakers should be associated with a more narrow range of values than utterances produced by Chill ones. We call this **Hypothesis 1**.

Second, we ask how the interaction between persona-based information and pragmatic reasoning is affected by how comprehenders see themselves with respect to the salient social qualities embodied by the speaker. As extensively discussed in sociolinguistic theory, indexical associations are not a static reflection of a speaker's identity; rather, they emerge from a dynamic, intersubjective interpretation process, which is ultimately shaped by how interlocutors orient themselves towards the social categories salient in the context (see Irvine and Gal 2000; Eckert 2008; Campbell-Kibler 2011; Levon 2014; Acton and Potts 2014 i.a.). As a result, to gain a comprehensive understanding of how imprecision resolution is affected by persona information, one must also consider how this effect is modulated by comprehenders' own relation to the social features of the speaker. As to the nature of such a possible modulation, prior work on the relationship between social information and sound processing suggests two equally plausible hypotheses. One is that effects of social information on speech processing are especially strong when comprehenders *share* the salient properties of the speaker. This possibility is suggested by the general observation that listeners' perception tends to track production, and thus to be especially accurate when involving linguistic forms that are part of the listeners' own repertoire – a proposal that has been put forward by work on phonetic processing on different phenomena, focusing on different kinds of social information (Plichta and

Preston 2005; Hay et al. 2006; Sumner and Samuel 2009; Fridland and Kendall 2012). Applied to our case, this would suggest that the hypothesized persona effects on the interpretation of numeral expressions should be especially prominent for comprehenders who see themselves as embodying the same social profile as the speaker. We'll henceforth refer to this hypothesis as **Hypothesis 2A**.

But other previous lines of work indicate that the reverse interaction is possible as well. In particular, it has been suggested that listeners are especially prone to rely on social information in speech processing when they do not share the same social properties as the speaker – or at least, when they believe that they don't. For example, listeners from Michigan have been shown to correctly categorize the same diphthong as raised when they were told that the speaker was from Canada, in line with the dominant stereotype; but erroneously perceived it as centralized when they were told that the speaker was from Michigan – even though raised diphthongs are routinely produced by speakers from Michigan as well (Niedzielski 1999). In a similar vein, Wade (2020) showed that Non-Southerner listeners exhibit convergence towards speakers labeled as “Southern” but *not* using any actual Southern speech feature in their speech, while Southern listeners did not exhibit the same behavior, once again highlighting a greater propensity to recruit stereotypical knowledge in speech processing for listeners whose identity is not targeted by the stereotype in question (see Babel 2010; Walker and Campbell-Kibler 2015 for similar patterns involving convergence and imitation). Taken together, these findings suggest the alternative possibility, which we'll henceforth refer to as **Hypothesis 2B**: that the predicted effects of persona on the interpretation of numerals should be stronger for comprehenders who do *not* see themselves as sharing the same traits as the speaker.

We now proceed to test out hypothesis in an experiment deploying a picture-selection task.

3 Personae and (im)precision: a picture matching task

We explore these hypotheses in an experiment utilizing a picture selection task, which we label COVERED SCREEN TASK.

3.1 Methods & Design

Our stimuli utilized visually displayed dialogues, like those in the norming study reported above, with variants resulting from manipulations that crossed two factors in a 2×3 design. Our first manipulation varied the persona embodied by the displayed characters (Nerdy: Arthur and Rachel vs. Chill: Alex and Eva). Each set of characters was framed in a conversation prefaced by a brief context sentence; in the dialogue, the female character and the male character would respond based on information they accessed by looking at their phone, uttering a quantity expression in the form of a round number. The presentation was identical to the cartoons in Figure ??, with the question and the response provided as text in the speech bubble.

After seeing the dialogue between the characters, participants were asked the question “Which phone is Arthur/Alex looking at” and were shown 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). Our second manipulation varied the fit between the visible screen and the one uttered by the character, with three levels: *Match* with displayed numbers matching the utterance²

Figure 4. Screen Fit Manipulation

Participants were instructed to select the visible screen ‘if you think that the information on the screen fits what is being said’ and to select the covered screen if they believed it wasn’t. The Match and the Mismatch conditions serve as controls and are expected to evoke **COVERED** responses at floor and ceiling levels respectively. By contrast, responses in the Imprecise condition crucially depend on the (im)precision standard 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 extension of the predicate – should lead to a **COVERED** screen response; and a relatively lenient 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. This means that **COVERED** choices in the Imprecise condition can be seen as an indicator of the

standard of precision used.

3.2 Materials

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

3.3 Procedure

The study was implemented and administered online on the PCIBex platform (<https://www.pcibex.net>). After providing informed consent, participants were shown the instructions. Participants entered their responses by pressing the key matching the letter displayed under the picture on the keyboard. Figure 5 provides a full illustration of the display that participants would see once all the elements appeared.

Figure 5. Display before making the choice (Condition: Nerdy, Imprecise)

Following the last trial, all participants, regardless of whether they had been assigned to the Nerdy or the Chill condition, were asked to complete a two question exit questionnaire aimed at assessing the degree to which they saw themselves as sharing the salient qualities of the speakers in the experiment. Participants responded by selecting a value on a 1-10 scale, with 1 indicating the minimum value and 10 indicating the maximum value. The two questions were presented incrementally.

- (2) a. I'd describe myself as: 1=not chill at all; 10=very chill

- b. I'd describe myself as: 1=not nerdy at all; ... 10=very nerdy

The linking hypothesis behind these questions is that the responses should open a window onto the degree to which participants see themselves as sharing the same defining quality as the characters seen in the dialogue; this self-assessment could thus provide a measure to explore how the effect of Persona on imprecision resolution is modulated by respondents' similarity to the speaker persona (see Hypothesis 2A-B). Specifically, the two scales above provide two pieces of information: the degree to which respondents saw themselves as featuring the social quality distinctive of the persona they had just seen in the study, henceforth labeled the *similarity trait* – i.e., “nerdy” for participants seeing the nerdy speakers; and “chill” for participants seeing the chill speakers; and the degree to which they saw themselves as featuring a quality that wasn't central irrelevant to the persona that they had just seen, henceforth labelled *secondary trait* – i.e., “chill” for participants seeing the nerdy speakers; “nerdy” for participants seeing chill ones. We thus assume that the only informative rating with respect to the similarity between respondents and characters should be the similarity rating. This linking hypothesis was independently explored and validated.⁴

3.4 Participants

306 participants (Mean Age=30; self-identified gender: Female=164; Male=137; Non-binary=5) were recruited online from Prolific and compensated \$1.30 (\$8/hr). All participants declared to be native speakers of English and provided informed consent approved by our university's IRB.

3.5 Results

3.5.1 Task Validation and overall Persona effect

Our first step is to confirm that response patterns in the control conditions behave as expected, and to take a first pass at assessing the effect of Persona in the critical condition. Recall that our control conditions were set up to lead to ceiling and floor levels of covered box choice rates; only in the imprecise condition did we expect participants' responses to depend on how they reason

about imprecision, and to be affected by our manipulations. We present the proportions of covered box choice rates across conditions in Figure 6. As one can observe, there is a clear step-wise effect of Screen Fit Manipulation, with ceiling and floor-level response rates for the controls, and the imprecise condition in the middle, as expected. Second, while no difference seems to be observable between the Nerdy and Chill persona in the Match and Mismatch condition, a contrast between Nerdy and Chill speaker is observable in the Imprecise condition, with the rate of COVERED choices higher for Nerdy speakers.

Figure 6. COVERED choice rate by Screen Fit and Persona

To confirm these patterns statistically, we fit a mixed-effects logistic regression with Screen Fit and Persona and their interaction as fixed effects, and by-Subject and by-Item random intercepts (the maximally complex random effect structure that would converge). To explore the intended contrasts, we extracted two sets of comparisons from the model using the *emmeans* package. First, we compared the rate of COVERED responses in the Imprecise condition with those in Mismatch and Match, collapsing across Persona levels. The pattern suggested by the graph was confirmed: the rate of COVERED choices in the Imprecise condition differed significantly from both the rate in the Match ($\beta=-5.76$; $SE=0.26$; $p < 0.0001$) and the Mismatch ($\beta=7.21$; $SE=0.35$; $p < 0.0001$). Second, we compared responses for Chill vs. Nerdy speakers within each Screen Fit level.⁵ As predicted, no difference between responses with the Nerdy vs. Chill persona was observed in the two control conditions (Match: $\beta= 0.65$; $SE=0.54$; $p = 0.22$; Mismatch: $\beta=0.84$; $SE=0.71$; $p = 0.23$). By contrast, a significant effect was found for the Imprecise condition, with Nerdy speakers associated with a higher rate of COVERED responses ($\beta=0.78$; $SE=0.26$; $p = 0.003$).⁶

3.5.2 The modulation of participants self-ascribed ratings

To zero in on how the effect of speaker persona was modulated by participants' self-ascribed qualities, we now focus on the Imprecise condition only, and explore the relation between Persona manipulation and participants' self-ascribed ratings as interacting predictors. Drawing on the assumption that self-ascribed nerdiness and self-ascribed chillness, i.e., what we called the *similarity*

traits above, provide a window into the degree to which respondents see themselves as sharing the distinctive quality of the Nerdy vs. Chill persona respectively, we pooled together these traits to create a *Similarity Index*, which we treat in our analysis as an independent variable tracking the degree to which participants saw themselves as sharing the same qualities of the speakers. Figure ?? illustrates the proportion of COVERED choices, by speaker persona, in the Imprecise condition in relation to the speaker persona and the Similarity Index. As one can see, the rate of COVERED choices appears to be higher for Nerdy speakers than for Chill speakers when the Similarity Index is lowest; and to decrease, and eventually disappear, as the Similarity Index increases.

Figure 7. COVERED responses by speaker Persona and Similarity Index.

To explore this pattern statistically, we fit a mixed-effect logistic regression with screen choice as our dependent variable, speaker persona and Similarity Index as interacting predictors, and random intercepts for Subjects and Items (the maximal random effect structure that would converge). Persona was a sum-coded categorical predictor, and Predominant Trait was a centered continuous predictor (i.e., the scale from 1-10 was transformed to range from -4.5 to 4.5). The output from the model is reported in Table 1.⁷

Consistent with the results from the first step of our analysis, we find a significant main effect of Persona ($\beta=-1.50$; $SE=0.37$; $z=3.91$; $p<0.0001$) with higher COVERED screen rates for Nerds. But this effect is dominated by a significant interaction ($\beta=0.42$; $SE=0.14$; $z=0.44$; $p=0.002$), such that the Persona effect is very strong when participants' similarity index is high, and disappears when it is low. To further understand the nature of the interaction, we re-ran the same model using treatment coding with Nerd and Chill as respective baselines. This reveals a significant *decrease* of COVERED screen choice rates as the Similarity Index increases in the Nerdy speaker condition ($\beta = -0.24$, $SE=0.09$, $z=-2.60$, $p=0.009$); and a marginally significant *increase* of COVERED screen rates as the Similarity Index increases in in the Chill speaker condition ($\beta = 0.18$, $SE=0.10$, $z=-1.70$, $p=0.08$). This suggests that the observed interaction between speaker Persona and Similarity Index is primarily driven by participants seeing the Nerdy speakers; but that, to a certain extent, is also present for participants seeing the Chill speakers.

4 General Discussion

Our results support two conclusions. First, numerals uttered by speakers whose social persona embodies the qualities indexed by precise speech – i.e. the Nerdy ones – receive a more narrow interpretations than numerals uttered by speakers whose social persona embodies the qualities indexed by imprecise speech – i.e. Chill ones. Second, this effect is modulated by comprehenders’ own positioning with respect to this persona: persona-induced differences in interpretation are most pronounced for comprehenders who see themselves as *not* embodying the distinctive qualities of the speaker; and less pronounced for speakers who instead see themselves as sharing them.

We now turn to elaborate on the implications of these findings, with an emphasis on two areas: what they reveal about the interface between different domains of meaning; and how they speak to the role of social information in linguistic cognition more broadly.

4.1 Social vs. Descriptive meanings: revisiting the socio-semantics interface

First, our findings unveil a novel connection between the descriptive and the socio-indexical domains of signification; they thus further illuminate the mutual interaction between two activities that are key to linguistic interaction, but which have been by-and-large investigated on in separate scholarly domains: how interlocutors’ exchange propositional information, and thus zero in on state the world; and how interlocutors project, convey and perceive information pertaining to who they are – i.e., their identity, personae, and personality. In particular, the differences in interpretation associated with the two personae manipulated in the study suggest that axes of social differentiation between speakers, and they way in which they come to be indexed by linguistic and non-linguistic signs, *do* impact the outcome of interlocutors’ endeavor to exchange propositional information – and ultimately zero in on the state of the current world. As a result, pragmatic reasoning emerges as a *socially embedded* process: one that crucially hinges on how comprehenders perceive and relate to the social qualities embodied by the speaker, and thus cannot be fully grasped if one abstracts away from such qualities. This observation, in turn, allows us to take a further step in the endeavor of exploring the interface between the descriptive and socio-indexical domains of

meaning, with two contributions being especially important.

On an empirical level, our results align with, and crucially extend, the idea that social meanings are productively and systematically inferred from the semantic and pragmatic properties of speech (Acton 2019; Beltrama and Staum Casasanto 2017; Beltrama et al. 2022; Glass 2015; Jeong 2021; Thomas 2021; see §1 for further details). In fact, the emerging picture is one in which the link between social and descriptive meanings is *bi-directional*: social meanings, including persona-based ones, are not just carried by and inferrable from semantic and pragmatic features, but also serve as a resource for interlocutors to navigate indeterminacy at the semantic and pragmatic level, and ultimately compute what a description refers to in a given utterance context.

An important point to address, in this respect, revolves around how one should characterize the nature of the interaction between social indexicality and pragmatic reasoning observed in our study. The approach adopted here allows us to establish a link between *contrasts* in meaning interpretation (i.e., rates of precise interpretations) and *contrasts* at the level of personae (i.e., Nerdy vs. Chill); this outlook, as discussed in §2.2., is in line with the idea, central to sociolinguistic and anthropological theory, that personae do not exist in isolation, but are interpreted relative to a broader *system of distinctiveness* (Irvine 2001; Eckert 2019): they inherently evoke, and cannot therefore be fully divorced from, the opposition to possible other constructs that could have materialized in the context. In our conceptualization of the study, Nerdy and Chill can be thought of as an instance of such an opposition. At the same time, the issue remains open as to what specific qualities of the personae we manipulated, if any in particular, drive the observed effects on the resolution of (im)precision, with two questions emerging as especially prominent. One question concerns whether a similarly significant contrast in the outcome of pragmatic reasoning would have been observed by comparing numerals uttered by Nerdy vs. those uttered by speakers embodying a personae whose indexical link to imprecision might have been less strong than it was for the Chill – e.g., a “young professional” one, to cite an example of a persona that is sufficiently distinct from the Nerdy one, and received attention in the previous literature (D’Onofrio 2018). The other question revolves around whether the association between numerical precision and a Nerdy persona

extends across different subtypes of this persona, or is instead specific to some specific aspects of the incarnations evoked in our experiment. As most social meaning categories, personae are indeed not internally homogeneous: different manifestations of nerdiness and chillness exist, which can be seen as sharing a common denominator, and yet differ with respect to other crucial properties of the people embodying them. For example, not all nerds are necessarily pedantic or uptight; and not all chill people are necessarily laid back or likable (see Pratt 2022 for a recent take). In this perspective, we see minimal variations of our paradigm – e.g., testing the effect on numerals’ interpretation of different, and appropriately normed, incarnations of nerdiness and chillness, as well as of other persona constructs in opposition with either of these two – as a promising avenue of future work to shed light on what specific constellation of qualities are primarily impactful on comprehenders’ pragmatic reasoning; and thus better theorize the nature of the socio-indexical links that drive the observed impact of persona-based information on meaning interpretation.

On a theoretical level, our findings also carry a broader impact on the way in which meaning interpretation has been generally theorized, especially in the fields of semantics and pragmatics. In particular, our findings question the traditional separation between the propositional and the socio-interactive dimension of meaning that has characterized much research in the study of meaning making processes. To be sure, the central role of these dimensions in linguistic communication has long been acknowledged by scholars investigating language across different approaches (Ochs 1992; Duranti 1997; Eckert 2019 among many others), even in the tradition of semantics and pragmatics. In particular, while laying the foundations of a framework that would inspire much work in subsequent years, Grice did observe that “there are all sorts of other maxims (aesthetic, social or moral in character) [...] that are also normally observed by participants in talk exchanges [...] and could also generate inferences” (Grice 1975: 47). Yet, work in the wake of this tradition has put such other factors aside, seeing them as somewhat tangential to the core strategies whereby interlocutors engage in the exchange of propositional information. But this assumption crucially breaks apart once we consider the results from our study, highlighting the importance of developing a comprehensive enough pragmatic framework that treats interlocutors’ distinctive identities and

personality – and the wealth of cues whereby these are made salient in the conversational context – as a central component to the process.

As to how one might go about undertaking this challenge, we suggest that a particularly promising option comes from *Rational Speech Act* models – a framework that draws on Grice’s view of communication as a cooperative activity (Frank and Goodman 2012; Kao et al. 2014; Goodman and Frank 2016), and places an emphasis on the idea that interlocutors engaging in conversation are aiming at maximizing their *utility* vis-a-vis the goal that they are pursuing (e.g., informing each other about the state of the world). This framework is especially appealing in that it has already been productively deployed to capture both (im)precision with numerals (Kao et al. 2014; Aparicio 2017) and the construction and negotiation of social identity in connection with speech sounds (Burnett 2017, 2019), making two key insights available: the idea that a speaker’s choice to describe things more or less precisely stems from striking a balance between accuracy, granted by choosing a high level of precision, and simplicity, facilitated by choosing a lower level (see Kao et al. 2014); and the idea that social information is part and parcel of the message conveyed by an utterance, to the point that interlocutors are not just pursuing the goal of describing reality, but also of projecting a certain persona successfully in interaction (Burnett 2017, 2019). Accordingly, one could recast participants’ response behavior as stemming from how they expect a speaker embodying a particular persona to weight accuracy versus effort: in particular, Nerdy speakers might be expected to put a premium on accuracy at the cost of incurring a higher effort; and Chill ones to prioritize effortlessness, at the cost of sacrificing accuracy. This reasoning, in turn, will lead comprehenders to adjust their interpretive behavior when it comes to determining the referent of the numeral – a pattern that could be captured by implementing in the model different expectations about the accuracy/complexity tradeoff for speakers embodying different personae, and allowing these expectations to impact comprehenders’ reasoning accordingly.

4.2 Indexicality and sociolinguistic cognition: the role of speakers and comprehenders

The second major takeaway of our results is that the impact of socio-indexical information on language processing extends beyond the phonetic and syntactic domain – i.e., those on which previous work primarily focused. This, in turn, presents the possibility of considering how our findings speak to the effects of social indexicality explored in other, seemingly distinct domains of sociolinguistic cognition.

On a general level, our findings align with those from previous work showing that comprehenders keep track of different levels of speaker information – including both demographic features (e.g., Strand 1999; Niedzielski 1999; Hay et al. 2006; Hay 2009; Babel 2012; Drager 2015; Staum Casasanto 2008; Sumner et al. 2014) and personae (D’Onofrio 2015, 2018, 2020) – across different domains of speech production and perception – e.g., phonetic categorization, imitation, and convergence (see §1). But on a more specific level, our findings further indicate that the effect of speaker persona on imprecision resolution crucially interacts with comprehenders’ own position with respect to these personae – and in particular, with whether comprehenders see themselves as sharing the social qualities also embodied by the speaker they are seeing in the study.

When it comes to considering how this pattern relates to those emerging from previous work, we have suggested that our findings support what we have referred to as **Hypothesis 2B**, motivated by findings from the sociophonetics literature showing that social information cued via stereotypical links between forms and social categories exclusively affected participants’ responses when they did not see themselves as targeted by the stereotype at stake (Niedzielski 1999; Babel 2010; Walker and Campbell-Kibler 2015; Wade 2020; see §2.3 for details). They instead speak against Hypothesis 2A, which predicted that comprehenders embodying a particular persona have fine-grained knowledge about generalizations and expectations related to this persona’s linguistic behavior, and should therefore be especially likely to recruit this knowledge when processing utterances by speakers similar to them.

As to what specific motivation drives the observed interplay between speakers’ and comprehenders’ features, we would like to highlight two prominent possibilities, both of which have been

proposed in the previous literature and are consistent with our results. The first is that participants identifying with the speaker (consciously or unconsciously) selectively block the indexical association from affecting their response choices, most plausibly to fend off potential negative evaluations of their shared persona. In contrast, when identification with the speaker is low, participants' own identity is not at stake, leading participants to be more susceptible to letting these indexical associations impact their behavior. This line of thought would be consistent with patterns from the previous sociophonetics literature, such as the observation that participants can resist, or altogether inhibit, convergence towards the interlocutor when it comes to linguistic features that are embedded in stigmatized stereotypes (Babel 2010; Walker and Campbell-Kibler 2015); that whether a particular social stereotype about a linguistic feature is endorsed vs. opposed by participants has a significant effect on the way in which they process this feature in an experimental task (Levon 2014); and that, as suggested by Niedzielski (1999), speakers are more comfortable ascribing stereotypically rich social characterizations to others, while portraying themselves as unmarked language users, both from a linguistic and a social perspective. Alternatively, the interaction between speaker persona and participants' identification could reflect different degrees of having to rely on stereotypes, as suggested by Wade (2020) to capture the different propensity of Southerners and Non-Southerners to converge towards Southern-labeled speech that lacked actual Southern features (see §2.3). Applied to our case, it's possible that the socio-indexical association between personae and precision tracks real-world patterns of numeral usage at best loosely, and ultimately inaccurately – as is indeed the case for most ideological representations of speech (see Gal and Irvine 2019 for further discussion). On this view, respondents sharing the same traits as the speaker may have more and richer experiences with the relevant type of speaker, making socio-indexical information less impactful on their response behavior; and respondents who share the same speaker's traits to a lesser extent may have less access to information drawing on actual usage, and thus rely on indexical associations more heavily.

Finally, we observe that a further source of information that might be affecting participants' responses in this context is the social persona of the addressees in the presented dialogues, i.e.

Rachel and Eva, who ask the question that the speakers answer using a numeral. In our study, the two interlocutors differed in gender but shared social characteristics, creating dyads that were homogeneous in terms of persona but not in terms of their demographic profile. While this was kept constant across the two speaker identity conditions, an obvious question is whether, and how, different combinations between these properties would have affected respondents' behavior. Especially intriguing, in terms of future research, is which of two possible alternative behaviors are observed if the dyads contained mismatching personae, e.g., with a nerdy speaker addressing a chill addressee (or vice versa). One alternative would be that participants posit some sort of accommodation or convergence to take place – i.e., that nerdy speakers are taken to describe quantities more similarly to how chill ones do, and vice-versa, resulting in mitigation – or possibly neutralization – of our persona effects. The opposite possibility would be that a conversation between divergent personae might be construed as leading to strengthened patterns of identity differentiation on the part of the speaker, translating into an amplified speaker persona effect. Taken together, we think that these considerations set up a promising hypothesis space that should be explored in future work to illuminate the cognitive processes involved in recruiting, and reasoning about, socio-indexical information in different domains of language comprehension.

5 Conclusion

We presented experimental data showing that the social persona embodied by the speaker, in interaction with participants' own self-ascribed characteristics, impacts the pragmatic reasoning leading to ascribing an interpretation to numerical expressions. These findings open up novel perspectives on the study of meaning at the interface of semantics, pragmatics and sociolinguistics, highlighting the importance of developing a comprehensive approach to the study of these dimensions of signification.

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 sociolinguistics of demonstratives. *Journal of Sociolinguistics* 18(1), 3–31.
- 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.
- Babel, M. (2010). Dialect divergence and convergence in new zealand english. *Language in Society* 39(4), 437–456.
- 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.
- Benor, S. B. (2004). ‘talmid chachams and tseydeykeses’: 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.
- 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. (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.
- 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).
- Drager, K. K. (2015). *Linguistic variation, identity construction and cognition*. Language Science Press.
- Duranti, A. (1997). *Meaning in linguistic forms*, pp. 162–213. Cambridge Textbooks in Linguistics. Cambridge University 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.
- Eckert, P. (2019). The individual in the semiotic landscape. *Glossa* 4(1)(14).
- Fairchild, S. and A. Papafragou (2018). Sins of omission are more likely to be forgiven in non-native speakers. *Cognition* 181, 80–92.
- Frank, M. C. and N. D. Goodman (2012). Predicting pragmatic reasoning in language games. *Science* 336(6084), 998–998.
- Fridland, V. and T. Kendall (2012, may). Exploring the relationship between production and perception in the mid front vowels of u.s. english. *Lingua* 122(7), 779–793.
- Gal, S. and J. T. Irvine (2019). *Signs of Difference: Language and Ideology in Social Life*. Cambridge University Press.
- 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.

- 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.
- 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.
- 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.
- 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.
- Irvine, J. and S. Gal (2000). Language ideology and linguistic differentiation. In *Regimes of language*, pp. 35–83. Santa Fe, NM: School of American Research Press.
- Jeong, S. (2021). Deriving politeness from an extended lewisian model: The case of rising declaratives. *Journal of Pragmatics* 177, 183–207.
- 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.
- Kiesling, S. F. (2018). Masculine stances and the linguistics of affect: on masculine ease. *NORMA* 13(3-4), 191–212.
- King, S. (2021). Rethinking race and place: The role of persona in sound change reversal. *Journal of Sociolinguistics* 25(2), 159–178.

- 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.
- Levon, E. (2014). Categories, stereotypes, and the linguistic perception of sexuality. *Language in Society* 43(5), 539–566.
- 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.
- Mazzarella, D., E. Trouche, H. Mercier, and I. Noveck (2018). Believing what you're told: Politeness and scalar inferences. *Frontiers in Psychology* 9, 908.
- Niedzielski, N. (1999). The effect of social information on the perception of sociolinguistic variables. *Journal of Social Psychology (Special Edition)* 18:1, 62–85.
- 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.
- Pinkal, M. (1995). *Logic and lexicon*. Dordrecht.
- Plichta, B. and D. R. Preston (2005, jan). The perception of /ay/ as a north-south stereotype in united states english. *Acta Linguistica Hafniensia* 37(1), 107–130.
- 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.

- Roberts, C. (1996/2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5(6), 1–69.
- 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.
- Staum Casasanto, L. (2008). *Experimental Investigations of Sociolinguistic Knowledge*. Ph. D. thesis, Stanford University.
- 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.
- Sumner, M. and A. G. Samuel (2009, may). The effect of experience on the perception and representation of dialect variants. *Journal of Memory and Language* 60(4), 487–501.
- 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. (2020). *The Linguistic and the Social Intertwined: Linguistic Convergence toward Southern Speech*. Ph. D. thesis, University of Pennsylvania.
- Walker, A. and K. Campbell-Kibler (2015, may). Repeat what after whom? exploring variable selectivity in a cross-dialectal shadowing task. *Frontiers in Psychology* 6.

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.

Endnotes

¹This idea is crucially reflected in the idea that these inferences must be calculable (Grice 1975): *ceteris paribus*, they should be reconstructed in similar ways by cooperative interlocutors, regardless of who they are.

²The Match condition came into two slightly different variants, counterbalanced across lists: one in which the visible number had .00 (e.g.: Uttered number: “\$200”; Visible screen number: “\$200.00”) for the decimal part; one in which it had non-zero figures (e.g. “\$200.17”). This variation was introduced to ensure that participants wouldn’t reason to the effect of rejecting any occurrence of visible screen that didn’t perfectly match the uttered number. A post-hoc analysis indicates that this variation didn’t affect the overall pattern of results (see Endnote 7)

³Links to full instructions and stimuli list available in the Supporting Information file.

⁴See Supporting Information file for details.

⁵Due to the challenge of interpreting interactions involving response rates in logistic regressions that closely approximate 0% and 100%, such as those for the Match and Mismatch condition, we chose not to directly interpret the interaction between Screen Fit and Persona.

⁶Separate regressions were fit on variants of the experiment corresponding to different Match condition implementations (see Endnote 3). Significant Persona effects were found in both models, suggesting that this variation didn’t impact the key pattern of results (.00: $\beta = 0.94$, $SE=0.43$, $p=0.03$); non-zero: $\beta = 0.81$, $SE=0.24$, $p=0.01$).

⁷The interaction Persona*Similarity.Index also obtains in a model that additionally includes the other trait (self-ascribed chillness for participants seeing Nerdy and self-ascribed nerdiness for participants seeing Chill characters): $\beta = 0.41$, $SE=0.14$, $p=0.03$.