(Im)precise personae: The effect of socio-indexical information 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 contribute to building a more integrative perspective between the socio-indexical and the referential domain of signification, highlighting comprehenders' social perception of the speaker as a crucial element informing pragmatic reasoning, and meaning interpretation more broadly.

Keywords: social meaning; personae; pragmatic reasoning; precision; numerals.
Word count: 9994

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, such as their demographic background, ideological orientation, and personality traits (Labov 1966; Ochs 1992; Silverstein 2003; Eckert 2008 i.a.). Both these dimensions of signification play a pivotal role in linguistic interaction, enabling conver-

sational participants to pursue two activities central to communication: exchanging propositional information, and thus zeroing in on the way things are in the world; and projecting and perceiving their identity and personality, thus reflecting and reshaping the social context in which interaction takes place. Accordingly, scholars across linguistic subfields have suggested that a comprehensive understanding of how meanings are produced, negotiated and ascribed in communication requires understanding the interaction between these domains of signification – that is, how interlocutors triangulate between them to signal and interpret legible content (Gumperz 1982; Silverstein 1985; Ochs 1992; Duranti 1997; Eckert 2019a), and how socio-indexical meaning can be understood "in the larger context of language as a system of meaning" (Eckert 2019b: 752).

Building on these insights, the question of how descriptive and socio-indexical meaning relate to one another has been addressed from two complementary angles. One revolves around how social meanings can be inferred on the basis of the semantic content of linguistic forms – a perspective recently developed by a growing line of work across sociolinguistics and pragmatics. In a paradigmatic study, Acton and Potts (2014) argue that the social meaning of solidarity and reciprocal affiliation conveyed by demonstratives *this* and *that* can be derived from the core semantics of these expressions – specifically, the presumption that the addressee can access the referent of the embedded noun phrase by considering the speaker's relation to entities in the discourse context. Similar inferential patterns from the semantic to the social plane have been unveiled for many other phenomena, including regular determiners (Acton 2019; Hunt and Acton 2022), intensifiers (Beltrama and Staum Casasanto 2017); modals (Glass 2015); discourse particles (Thomas 2021); and types of speech acts (Jeong 2021; see Beltrama 2020 for an overview). A crucial contribution of this work has been to show that social meanings, when conveyed by expressions above the sound level, can be grounded in, and shaped by, these expressions' semantic and pragmatic properties, highlighting a principled connection between different planes of signification.[1]

The other angle on the connection between descriptive and social meanings revolves around the reverse question: how does socio-indexical information, conversely, shape comprehenders' interpretation of descriptive content? While this perspective drew less attention than the other

in recent work, it originally received a strong impulse from work in interactional sociolinguistics, which advocated the view that the understanding of referential meaning, far from being a simple act of decoding, is a dynamic enterprise in which conversational participants are required to activelty fill the gaps left open by the literal meaning of words and sentences. Central to this process, in particular, is what Gumperz refers to as *conversational inferences* (Gumperz 1977, 1978, 1982), through which conversational participants "arrive at situated or context-bound interpretations of what is intended at any one point in an interaction" (Gumperz 1978: 395) by integrating different layers of contextually salient information – ranging from the phonetic properties of the utterance, to the stylistic and identity categories salient in the conversational context. For example, a bilingual mother addressing referentially identical expressions at her kid in different languages (e.g., "come here" in English vs. "ven acá" in Spanish) might be taken to be performing radically different speech acts depending on the socio-indexical value that the addressee ascribes to her linguistic choice in the context – e.g., issuing a warning with English vs. making a personal appeal to the kid's cooperation with Spanish (Gumperz 1977).

The amenability of referential meaning to being shaped by socio-indexical information, in turn, raises an important question. What generalizations can be drawn about the systematic reasoning patterns whereby interlocutors recruit social meaning in the interpretation process? More specifically: what regularities can we observe in the interpretive strategies whereby comprehenders utilize socio-indexical information to draw inferences aiming at the descriptive content of an utterance? The ethnographic approach of interactional sociolinguistics shed considerable light on the situation-specific dynamics at play in the peculiar context under investigation; yet, the strongly local orientation of this work was less equipped to illuminate the systematic principles that shape how different comprehenders recruit information about the social context to interpret descriptive meaning. Notably, this limitation was acknowledged by scholars Gumperz himself, as he observed that "it might be argued that conversational inference varies from person to person and is too situation-specific to be analyzed in general terms" (Gumperz 1978: 408).

At the same time, the issue around the systematicity and generalizability of the effects of

socio-indexical information on pragmatic inferences similarly escaped the scope of more global, generalization-oriented approaches to meaning interpretation – particularly those in formal pragmatics and analytic philosophy. These traditions notably focused on unveiling the universal conversational principles that underlie communicative exchanges – and how comprehenders uniformly reason about them to draw inferences (Searle 1969; Grice 1975; Horn 1984; Sperber and Wilson 1995; Levinson 2000 among many others). For instance, interlocutors across conversational contexts have been shown to systematically take quantifiers such as *some* as having an upper-bounded interpretation (i.e., as meaning "some, but not all") – an inference generated by the general assumption that the speaker, by obeying the Quantity Maxim, would have used a stronger term like *all* if they had been in the position of doing so truthfully (see Grice 1975; Horn 1984). Yet, within this line of work, little has been said about the role played by social meaning in the interpretive process. In fact, the move of abstracting away from the the socio-indexical dimension of signification is especially evident in Grice's (1975) foundational and highly influential theory of meaning, in which comprehenders are treated as socially undifferentiated interlocutors – i.e., idealized and cooperative rational agents whose axes of social differentiation are irrelevant to how descriptive meanings are interpreted; and meaning-related inferences are framed as *calculable* – i.e., amenable to being reconstructed in similar ways by different interlocutors, regardless of who they are, once other general contextual factors are kept constant.

The emerging picture is one in which much remains to be seen on the interplay between two key elements of the meaning making process: the wealth of socio-indexical information contextually available to interlocutors; and the reasoning dynamics whereby comprehenders assign a descriptive meaning to linguistic utterances. In particular, while work in sociolinguistics and discourse analysis illuminated the socially embedded nature of many conversational inferences, it didn't explore the extent to which these inferences can generalize across different contexts and/or comprehenders; conversely, while work in pragmatics and philosophy did focus on the systematicity and generalizability of interpretive processes, it hardly considered axes of social differentiation between interlocutors as relevant to this domain. Accordingly, we see the endeavor of bridging

this gap as a crucial step towards furthering the study of meaning making and interpretation – one that could integrate insights from different scholarly traditions to better understand how different domains of signification enable interlocutors to exchange content in communication.

It is important to observe that, in addition to these considerations, the task of undertaking this enterprise is further motivated on an empirical level. In particular, a rich body of experimental work in sociophonetics has provided evidence that social meanings serve as a cue for comprehenders to parse linguistic input across different domains of sound processing – e.g., sound categorization and recognition; convergence; imitation. Systematic 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 2010; Drager 2015; Staum Casasanto 2008; Sumner et al. 2014; Wade 2020); but also to 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). 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). Combined with the finding discussed above that social meanings can be grounded in the semantic and pragmatic properties of linguistic properties, these results suggest that socio-indexical meanings consistently shape other domains of language understanding such as phonetic processing, further grounding the hypothesis that socio-indexical information could similarly shape comprehenders' reasoning in other domains of linguistic cognition – and specifically, meaning interpretation – to a similar degree of systematicity.

In this paper, we set out to explore this hypothesis through the lens of the phenomenon of *(im)precision* in the interpretation of numerals, which we now turn to introduce.

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 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 variaton 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 effortlessness, 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 1-2.

Figure 1: Arthur and Rachel

Figure 2: Alex and Eva

To ascertain the viability of the persona contrast implemebted above, we conducted a norming study comparing the social evaluation of the two sets of characters in Figure 1-2. 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 3-4, 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.

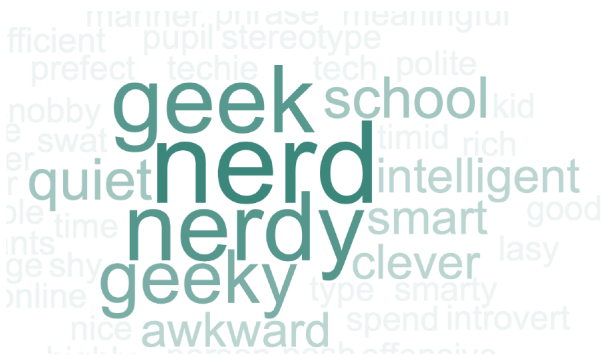

Figure 3: Qualities ascribed to Arthur.

Figure 4: Qualities ascribed to Alex

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 5, 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.

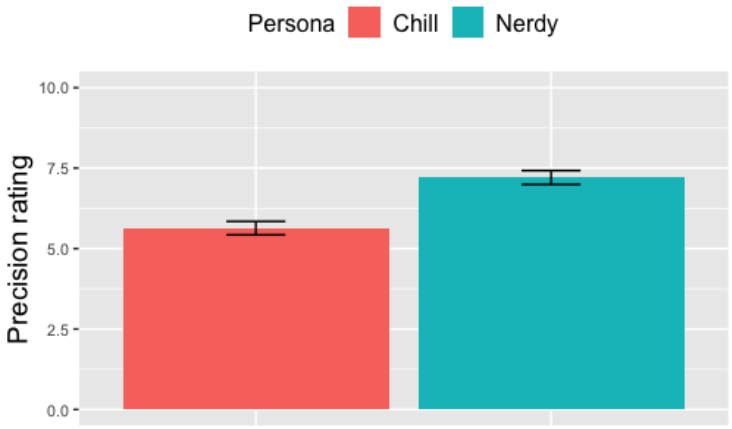

Figure 5: Expected precision for Chill vs/ Nerdy characters

2.3 Personae and (im)precision resolution: hypotheses

We can now utilize this contrast to explore how comprehenders' perception of the speaker's persona personae 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 embody 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 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**.[2]

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 2, 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 utterance[3]; *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 difference between the uttered and displayed numbers here varied between 5% and 18% of the first digit's unit.

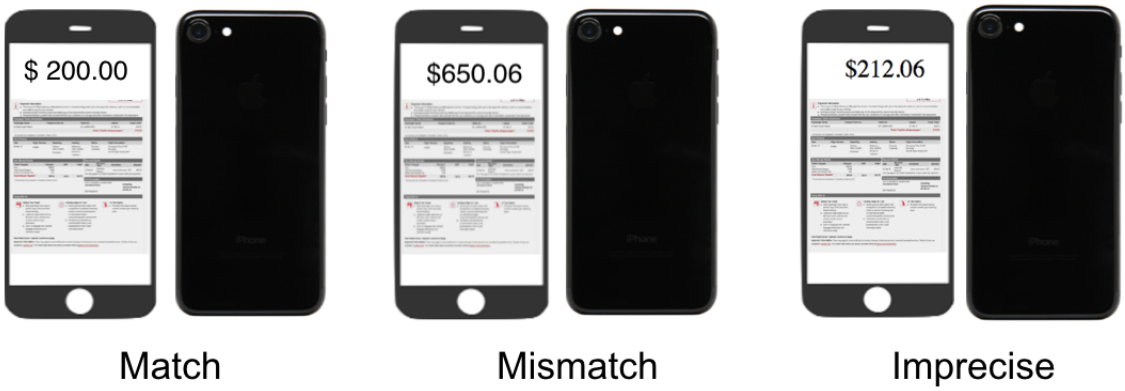

Figure 6: Screen Fit Manipulation

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.[4]

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 7 provides a full illustration of the display that participants would see once all the elements appeared.

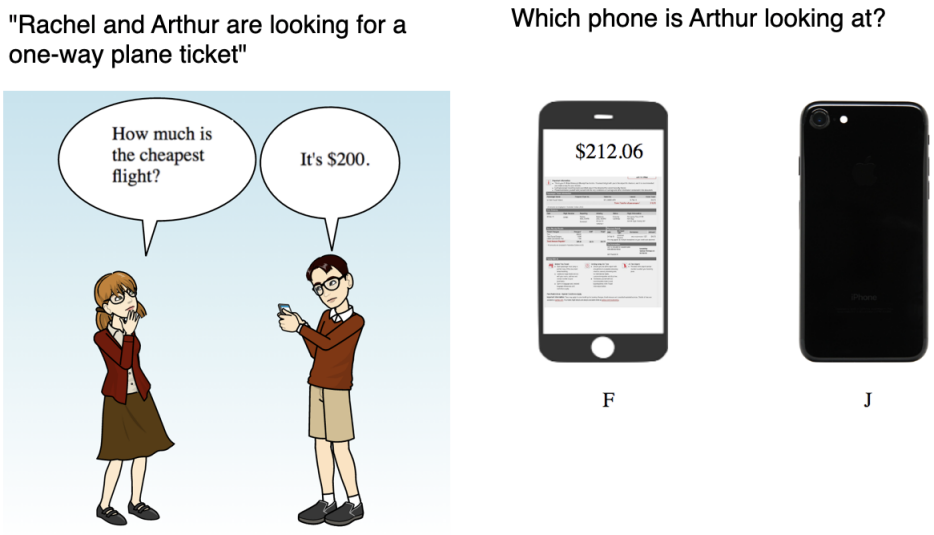

Figure 7: Display before making the choice (Condition: Nerdy, Imprecise)

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 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.[5]

3.4 Participants

306 participants were recruited online from Prolific and compensated $1.30 ($8/hr; DEMOGRAPHICS). 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 8. 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.

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

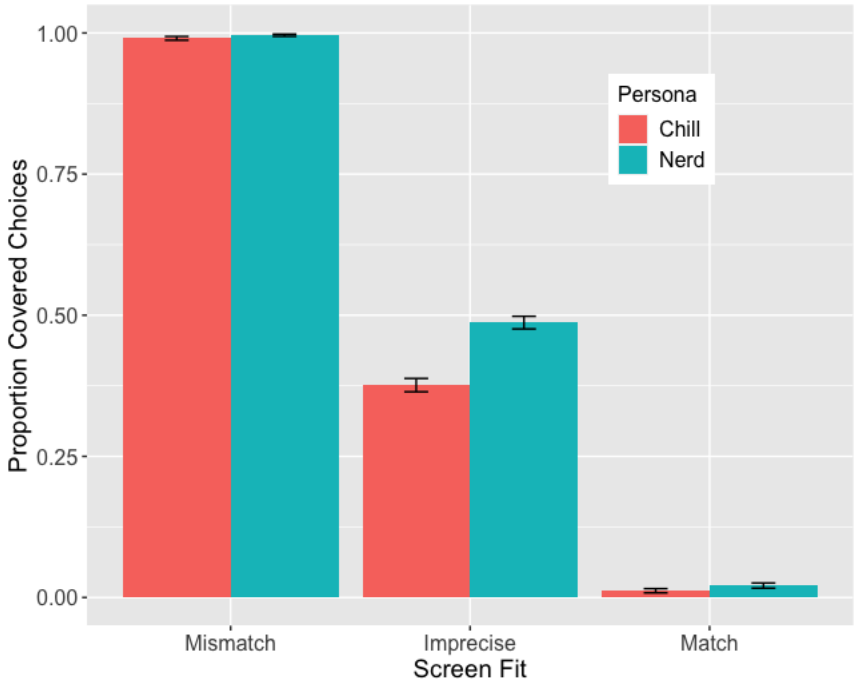

Figure 8: COVERED choices across Screen Fit, split by Persona

(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 (β=-5.76; SE=0.26; p <0.0001) and the Mismatch (β=7.21; SE=0.35; p <0.0001). Second, we compared responses for Chill vs. Nerdy speakers within each Screen Fit level.[6] As predicted, no difference between responses with the Nerdy vs. Chill persona was observed in the two control conditions (Match: β= 0.65; SE=0.54; p = 0.22; Mismatch: β=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 (β=0.78; SE=0.26; p = 0.003).

3.6 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 9 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.

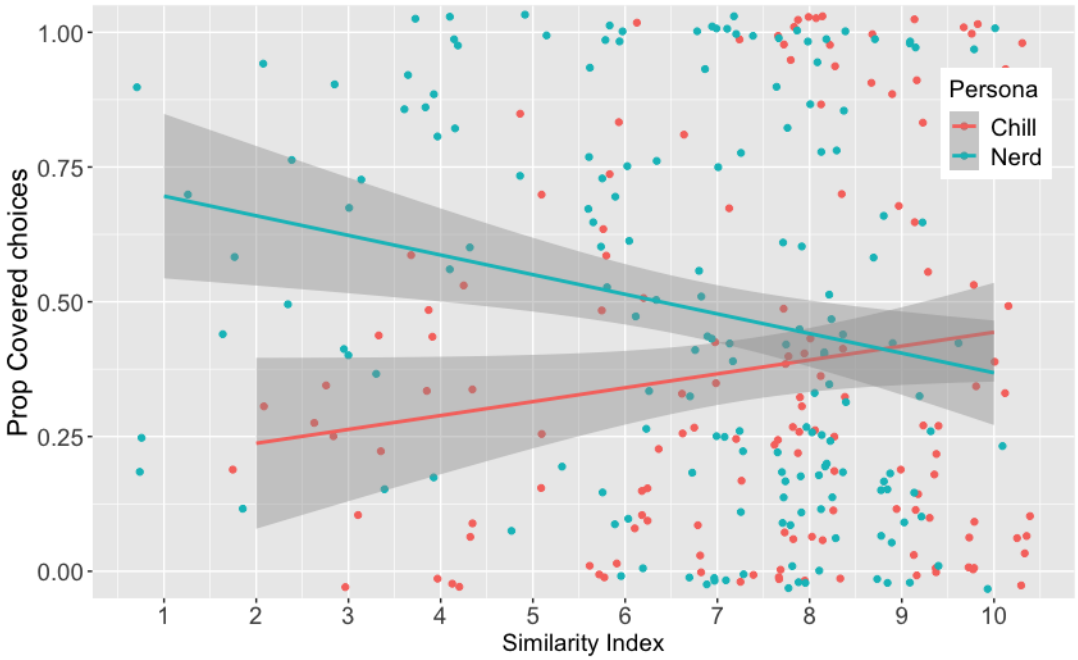

Figure 9: COVERED responses by speaker Persona and Similarity Index

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 (β=-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 (β=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 socio-indexical information in linguistic cognition.

4.1 Meaning across dimensions: enriching the socio-semantics interface

The first implication of our findings is that axes of social differentiation between speakers, and the way in which they come to be indexed by linguistic and non-linguistic signs, turn out to shape the general patterns of pragmatic reasoning whereby interlocutors make sense of an utterance; they thus afford a novel perspective on how different layers of signification jointly inform the production and interpretation of content in communication. As discussed in §1, exploring how different dimensions of signification allow interlocutors to signal and interpret legible meanings has long been a central goal in the study of human language (Gumperz 1977; Silverstein 1985; Ochs 1992; Duranti 1997; Eckert 2019a,b). This enterprise has been revived and expanded by a recent line of work at the interface of pragmatics and sociolinguistics, which showed that social meanings are productively and systematically inferred from the semantic properties of speech across a variety of linguistic expressions (Acton and Potts 2014; Acton 2019; Beltrama and Staum Casasanto 2017; Beltrama et al. 2022; Glass 2015; Jeong 2021; Thomas 2021; Hunt and Acton 2022; see §1 for further details), motivating the hypothesis that socio-indexical information can also systematically guide inferences targeting semantic interpretation.

Our results crucially provide evidence supporting this hypothesis, highlighting two distinctive properties of the interpretive inferences targeting numerical precision. First, they are *socially embedded*: that is, they hinge 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. Second, they are *systematic*: that is, they are drawn by different comprehenders in a consistent way, which generalizes beyond the here-and-now of the single interaction. This suggests that the way social meanings are integrated in semantic interpretation, while intimately tied to the specifics of the conversational setting, is similar across comprehenders, allowing us to bring together the

outlook of distinct scholarly traditions concerned with the study of meaning.

To begin with, the effect of speaker persona on precision inferences sheds a new light on the perspective on conversational inferences developed by interactional sociolinguistics. Scholars in this tradition already highlighted how everyday language relies on simultaneously conveyed referential and indexical signs, "one working in conjunction with the other" (Gumperz and Cook-Gumperz 2007); and suggested that the interpretation of descriptive meaning cannot be conceived of independently of the social context – a position largely shared by subsequent work across discourse analysis and intercultural pragmatics (Cook-Gumperz 1992; Gumperz and Cook-Gumperz 2007; Tannen 2008; Terkourafi 2008; see Gordon and Tannen 2021 for an overview). Our results suggest that, at least to a certain extent, the highly localized dynamics behind the integration of contextualization cues coalesce into consistent patterns of pragmatic reasoning. It follows that, even though they are deeply grounded in the specificity of a the contextual setting (see §1), conversational inferences are drawn consistently, and can thus be investigated, on a larger scale than the contingency of a single interactional setting – even when they are guided by the speaker's distinctive social characteristics, as opposed to the more abstract properties of the utterance context. Of course, this doesn't mean that the effect of persona on numeral interpretation should be taken to be universal across speakers, or invariant across contexts. In fact, the study of numerals' interpretation is inscribed in a specific cultural context – for example, that of a community of speakers of English that, while from different regions, presumably share a set of ideologies and norms about the use of quantity expressions in communication, and the link between descriptive precision and related concepts such as truth, falsity and accuracy. As such, the findings emerging from our study remain to a certain extent local and context-bound, much like any endeavor to characterize linguistic patterns and practices – including Gricean and other philosophical approaches aiming to characterize meaning interpretation via a finite set of maxims and principles (see Ochs 1976; Irvine and Gal 2000 for a critical discussion of these approaches).

At the same time, precisely the systematicity observed in the response patterns ultimately makes it possible to build a bridge between sociolinguistic and more abstract, reference-oriented

approaches to the study of meaning. In particular, our findings call into question the traditional exclusion of the socio-indexical information – and axes of differentation between speakers more broadly – from the purview of philosophical and semantic theories of meaning. They thus highlight the importance of developing a comprehensive enough pragmatic framework that, while maintaining its commitment to unveiling the principles and reasoning patterns in virtue of which pragmatic inferences are drawn, treats interlocutors' distinctive identities and personality – and the cues that index them in the conversational context – as a central component to the reasoning dynamics that underlie the resolution and interpretation of descriptive meaning.

An important point to address, in light of these observations, 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., Nerd 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 2019a): 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 e.g., Pratt 2021). In this perspective, minimal variations of our paradigm – e.g., testing the effect on numerals' interpretation of different incarnations of nerdiness and chillness, as well as of other persona constructs in opposition with either of these two – provide 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 impact of persona-based information on meaning interpretation.

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

The second major takeaway of our findings revolves around the link between social indexicality, personae and cognition. As discussed in §1, work in sociophonetics showed that comprehenders keep track of different levels of speaker information across different domains of speech production and perception – e.g., phonetic categorization, imitation, and convergence. These effects have been observed in relation to macro-level demographic features, such as gender, race and geographical origin (e.g., Strand 1999; Niedzielski 1999; Hay et al. 2006; Hay 2009; Babel 2012; Drager 2015; Staum Casasanto 2008; Sumner et al. 2014); but also, and crucially, with respect to more local persona-based constructs (D'Onofrio 2015, 2018, 2020). Our findings suggest that persona-based constructs are even more central language understanding than previously known, enriching prior work on sociolinguistic cognition in two different ways.

First, they expand the range of areas of language understanding that have been shown to be influenced by persona-based information beyond those of phonetic and morpho-syntactic processing. While the category of persona has gained an increasingly central role in theories of social meaning and sociolinguistic cognition, much remains to be discovered on how "personae are rep-

resented and connected in the mind, and how these connections influence processes of linguistic perception" (D'Onofrio 2020: 9). In this regard, processes of meaning interpretation emerge as a viable testbed to shed further light on the link between personae and linguistic cognition. In particular, while it has been shown that meaning resolution can be affected by speaker information (e.g., political affiliation: Mahler 2022; language proficiency: Fairchild and Papafragou 2018) and politeness considerations (Bonnefon and Villejoubert 2006; Zhang and Wu 2020 i.a.), its sensitivity to socio-indexical associations remains largely unexplored. Our study crucially highlight personae as yet another dimension of social signification that shapes this process, extending the range of realms of language processing known to be sensitive to persona-based constructs, and thus contributing to developing a more unified view of seemingly distinct domains of linguistic cognition across sounds and meaning.

Second, our findings 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. In particular, our results support what we have referred to as **Hypothesis 2B**, motivated by prior research 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 should be more likely to reason about this persona's characteristics when processing linguistic input (see §2.3 for details)

As to what specific motivation drives the observed interaction, two prominent possibilities are consistent with our results, both of which have been proposed in the previous literature. 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 – and in particular, the general pattern reported in the social psychology literature that individuals have more complex and varied representations of perceived in-group members, and more schematic, and ultimately inaccurate, representations of out-group members (see in particular Park and Judd 1990; Judd et al. 1991; Boldry et al. 2007 on the Outgroup Homogeneity Effect). An explanation along these lines is 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 incorrectly – 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. We believe that these considerations set up a promising hypothesis space to be explored in future work.

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 continuing to develop a comprehensive approach to the study of these dimensions of signification.

Endnotes

[1]Furthermore, recent work in pragmatics proposed that the mechanisms behind the inference of social meanings can be conceptualized and formalized in a similar way to those underlying the inference of descriptive meanings (see Burnett 2017, 2019; Acton 2022 for different approaches).

[2]This hypothesis amounts to positing that the production/perception parallelism in phonetic processing extends to pragmatic processing as well – a possibility that has indeed been put forward by recent work on models of pragmatic inference (Waldon and Degen 2020).

[3]The Match condition came into two slightly different variants: one in which the visible number was completely identical to the uttered one; and one in which it only matched in the integer, creating a minimal discrepancy (e.g.: Uttered number: "$300"; Visible screen number: "$300.17"). This variation was introduced to ensure that participants wouldn't reason to the effect of rejecting any occurrence of visible screen that doesn't perfectly match the uttered number. A post-hoc comparison revealed no effect between whether the Match condition was implemented as a fully perfect Match or a near-perfect Match. See Supporting Materials for details

[4]Links to full instructions, data, R code, and a complete list of stimuli (experimental + fillers) are available at: Supporting Materials.

[5]See Supporting Materials. for details

[6]A more direct approach would have been to directly interpret the interaction between Screen Fit and Persona. But due 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 follow this route.

[7]But the key pattern also obtains in a more complex model that includes secondary traits; see linked Supporting Materials.

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). Sociophonetics, semantics, and intention. *Journal of Linguistics 58*(3), 465–494.

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 tsedeykeses': Language, learnedness, and masculinity among orthodox jews. *Jewish Social Studies 11*(1), 147–170.

Boldry, J. G., L. Gaertner, and J. Quinn (2007). Measuring the measures. *Group Processes and Intergroup Relations 10*(2), 157–178.

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.

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.

Cook-Gumperz, J. (1992). Gendered contexts. In *The Contextualization of Language*, pp. 177. John Benjamins Publishing Company.

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. (2019a). The individual in the semiotic landscape. *Glossa 4(1)*(14).

Eckert, P. (2019b). The limits of meaning: Social indexicality, variation, and the cline of interiority. *Language 95*(4), 751–776.

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

Fridland, V. and T. Kendall (2012). 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).

Gordon, C. and D. Tannen (2021). *The Routledge Handbook of Language, Gender, and Sexuality*, Chapter Interactional sociolinguistics. Foundations, developments, and applications to language, gender, and sexuality, pp. 181–196. Routledge.

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.

Gumperz, J. J. (1977). The sociolinguistic significance of conversational code-switching. *RELC Journal 8*(2), 1–34.

Gumperz, J. J. (1978). Dialect and conversational inference in urban communication. *Language in Society 7*(3), 393–409.

Gumperz, J. J. (1982). *Socio-cultural knowledge in conversational inference*, pp. 153–171. Studies in Interactional Sociolinguistics. Cambridge University Press.

Gumperz, J. J. and J. Cook-Gumperz (2007). *Chapter 16. A postscript: Style and identity in interactional sociolinguistics*, pp. 477–502. Berlin, New York: De Gruyter Mouton.

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.

Judd, C. M., C. S. Ryan, and B. Park (1991). Accuracy in the judgment of in-group and out-group variability. *Journal of Personality and Social Psychology 61*(3), 366–379.

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.

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

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

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. (1976). The universality of conversational postulates. *Language in Society 5*(1), 67–80.

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.

Park, B. and C. M. Judd (1990). Measures and models of perceived group variability. *Journal of Personality and Social Psychology 59*(2), 173–191.

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

Plichta, B. and D. R. Preston (2005). 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.

Pratt, T. (2021). Affect in sociolinguistic style. *Language in Society 52*(1), 1–26.

Searle, J. (1969). *Speech acts*. Cambridge: Cambridge University Press.

Silverstein, M. (1985). Language and the culture of gender: At the intersection of structure, usage, and ideology. In E.Mertz and R. Parmentier (Eds.), *Semiotic Mediation*, pp. 219–259. Orlando: Elsevier.

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.

Sperber, D. and D. Wilson (1995). *Relevance: Communication and cognition*. Blackwell.

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). The effect of experience on the perception and representation of dialect variants. *Journal of Memory and Language 60*(4), 487–501.

Tannen, D. (2008). *Talking Voices*. Cambridge University Press.

Terkourafi, M. (2008). *Toward a universal notion of face for a universal notion of cooperation*, pp. 313–344. De Gruyter Mouton.

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.

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

Walker, A. and K. Campbell-Kibler (2015). Repeat what after whom? exploring variable selectivity in a cross-dialectal shadowing task. *Frontiers in Psychology 6*.

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.