Authority and Gendered Speech: Cantonese Particles and Sajiao

We investigate differences between masculine and feminine speech in modern Cantonese, a Chinese language mainly spoken in Hong Kong and the Guangdong province of China. With respect to the study of gendered speech, Cantonese is interesting on at least two accounts.

1. Cantonese is rich in Sentence Final Particles (SFPs). These elements are interesting from the perspective of gendered language for a number of reasons. Firstly, stereotypes about gendered speech in Cantonese usually involve considerations on SFPs; some SFPs are perceived by speakers as carrying gendered meaning or being biased toward a particular gender. Still, unlike many other markers of gendered speech, they are easily observable because they form a closed class or reasonable size. Finally, they convey expressive meaning, which has been shown to be linked to gender (e.g. in gender on pronominals and indexicals), and so form part of a larger category of expressive gender-related items.

2. Cantonese, along with at least Mandarin, recognizes a particular register of speech called *sajiao* (SJ) in Mandarin. It refers to a "babyish form of persuasion" usually associated with young women. This register is linguistically interesting for several reasons. First, it is easily identified and often discussed by speakers and thus is usually consciously used by speakers and meant to be recognized, unlike other registers. Second, SJ is indicated by various markers, among which is the use of specific SFPs.

In spite of these considerations, little has been written on the topic of Cantonese gendered speech, and no large scale corpus study of the differences between masculine and feminine speech has been reported. We thus propose to fill this gap. First, we observe the distribution of a large inventory of SFPs in Hong Kong Cantonese across genders. For the gender biased particles we provide an explanation by relating them to a difference in epistemic authority: while men are generally deemed to be reliable, the same does not hold for women who need to resort to different strategies to overcome these biases.[1]

We then open up with some considerations on the validity of using data from fiction (e.g. TV drama or movies) to study gendered speech. We conclude with a discussion on SJ and propose a game-theoretic explanation for its mechanisms.

1 Cantonese Sentence Final Particles

Cantonese has a large inventory of Sentence Final Particles (SFPs). Their number and defining features vary in the literature, but usually around 30 to 40 distinct particles are distinguished. These appear in utterance final position, often in clusters of two or more SFP (Kwok, 1984; Matthews and Yip, 2011). The meaning of these particles is also quite varied; they can convey information about the (assumed) epistemic status of the participants, change the speech act associated with the utterance, indicate some information about the mood of the speaker etc.

A commonly held belief about these particles is that women use more of them than men, that they are reserved for very informal occasions and that more educated speakers use them less often or barely at all. These beliefs are justified either by the idea that SFPs mostly convey the feelings of the speaker, and that feminine speakers are more prone to express these, or that SFPs are related to politeness issues, and that women need to show more politeness than men (Light, 1982; Erbaugh, 1985; Chan, 2000). Surprisingly, such claims have never been empirically assessed.

The most extensive work on the gendered use of Cantonese SFPs is summarized in (Chan, 2002) who reports results partially gathered in a corpus of dialogues taken from a Chinese mainland Cantonese TV drama or based on intuitions. She identifies the following particles as being extensively used by feminine speakers compared to masculine ones: *zek1, ho2, wo3, laa3, laa1, aa3* and *aa5*.

[1] Before moving on, a brief aside on terminology: we will refer to 'men,' 'women,' 'male speakers,' 'female speakers,' 'masculine speech,' and 'feminine speech;' all these terms should be understood as referring to speech styles associated stereotypically with the speech of individuals presenting (or aiming to present themselves) as male or female, or to such individuals. We are not endorsing these stereotypes or the ideologies associated with them. The reader is welcome to replace these terms with more nuanced ones in case of discomfort; our terms can be understood as simple abbreviations used for convenience and variety of prose.

1.1 Distribution of SFP in a Cantonese corpus

The corpus used by Chan was rather limited, and even though the dialogues in the show aimed at being as natural, the question of their authenticity remains open. To address the question of the distribution of Cantonese SFPs across genders, we therefore relied on the conversation section of the *Hong Kong Cantonese Corpus* (HKCanCorp), a PoS tagged corpus of modern Hong Kong Cantonese (Luke and Wong, 2015).

Based on lists of SFPs found in the literature, we compiled a list of 49 Chinese characters which could be used to indicate a SFP in the corpus (since SFPs are very colloquial, and Cantonese has a limited written tradition, different characters can usually be used to write SFP, e.g. the particle *laa1* can either be written 喇 or 啦; conversely the character 喎 can correspond to any of the particles *wo3, wo4* or *wo5*). We then extracted all utterances containing a character in that list, provided that either the character unambiguously corresponds to a SFP (e.g. 喎) or has been tagged with the appropriate PoS (the corpus has no specific PoS for SFP, but uses the related "語氣詞/Modal Particle" tag). Based on this list, we were able to calculate how many times each particle was used by a feminine or masculine speaker, along with the number of utterances with and without any of these SFPs were used. Table 1 summarizes the overall results.

	No SFP	At least one SFP
Feminine Utt.	1961	5413
Masculine Utt.	1593	3082

Table 1: Frequency of utterances with/without the use of a SFP in the *HKCanCorp*

According to a chi-squared test for independence, there is a statistically highly significant but weak correlation between gender and the use of at least one SFP by the speakers in the corpus ($\chi^2 = 77.006; df = 1; p < 0.001; \varphi = 0.079$). Feminine speakers seem to use SFP more frequently than men, as is usually claimed, but the claim that men seldom use them is not tenable since two thirds of their utterances involve at least one such element.

To investigate the usage of individual particles, we visualized the Pearson residuals of a fit of the contingency table of the frequencies of each SFP by gender. The `vcd` package (Zeileis et al., 2007) of `R` was used to produce the figure below, showing the results of the particles for which the residuals suggest a different usage between genders (along with some of the SFPs mentioned by Chan).

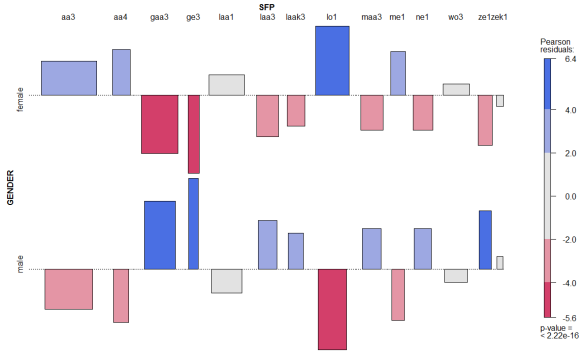

To further investigate the bias of the SFP which stand out in the above figure, Generalized Linear Mixed Models with optimized random structures were fitted in `R` to measure how well the relative frequency of usage of a particle predicts the gender of the speaker (treated as a binomial dependent factor). Recordings (indicated by the relevant file name) were treated as a random factor: they approximate the setting and formality of the conversation which might affect SFP usage. Model comparisons via likelihood ratio tests against a baseline model using the normalized utterance count as sole predictor and against a maximal model using all the baseline and all SFP as predictors were used to determine whether a given particle had a significant effect. Combining the results of this analysis with the association plot shown above, we summarize in Table 2 the SFP which are significantly associated with a gender.

Feminine Bias	Masculine Bias
aa3	*ge3*
laa1	*laa3 (laak3)*
lo1	*ze1*
me1	
wo3	

Table 2: SFP with a significant gender bias

1.2 Interpretation

A first observation is that out of the SFPs mentioned by Chan, the particles *ho2* and *zek1* were not found to have a significant gender bias. The

case of *zek1* will be discussed in Sec. 2; here we focus on the other SFPs. It is beyond the scope of this abstract to precisely describe the meaning of each particle mentioned in table 2. Here we stick to brief characterizations, mostly based on (Kwok, 1984) and (Matthews and Yip, 2011).

- *aa3* is the most common SFP (26.7% of the SFP tokens in the corpus). It is usually described as a "softener" of assertions and questions. It also involves a degree of evidentiality, it cannot be used if the speaker has no evidence for his claim.
- *laa1* marks invitations and mild commands.
- *lo1* indicates that its prejacent is expected and encourages agreement and cooperation. (Luke, 1990) describes it as a a marker used to negotiate successful conversation endings.
- *me1* is a polar question marker indicating that the speaker has a low belief in the affirmative.
- *wo3* indicates unexpectedness of either the content of its prejacent or the current conversational move (Hara and McCready, in press).
- The SFPs *ge3*, *laa3* and *laak3* are all described as markers of assertion, with different degrees of emphasis and finality. *Laak3* is supposed to be stronger than *laa3* in marking the relevance of their prejacent, and *ge3* marks the high confidence of the speaker in the truth of the prejacent.
- *ze1* downplays the importance of its prejacent.

We propose that the distribution of particles across genders is related to a difference in *epistemic authority*. (McCready and Winterstein, 2017) show using survey data that generally, feminine speakers appear to be assigned a low degree of credibility compared to masculine speakers outside of what is considered to be their domain of expertise. Their explanation is that women in general are assigned a lower degree of epistemic authority than men in Chinese society. Therefore, they cannot overtly refute other people's claims or "inflict" content on others. Men on the other hand can. This is reflected by their use of SFPs which either seek agreement or mark assertiveness.

1.3 Formalization

We present a formalization of this idea using the theory of pragmatic reliability proposed by (McCready, 2015) and extended by (McCready and Winterstein, 2017) to the case of gendered speech in particular. In this theory, an agent's judgement about the reliability of a speaker is arrived at on the basis of two factors. The first is the history of interaction between the agent and the speaker: the more communications of the speaker that properly convey accurate information, the higher the degree of reliability the agent will assign. Second, the baseline likelihood that the speaker is reliable at all will be assigned on the basis of heuristics about the kinds of agents that are reliable. Often, such judgements involve social factors: a professor will, for some domains, be deemed more reliable than a homeless person, for example.

It has been pointed out by (e.g.) (Fricker, 2009) that, in general, men have a higher degree of epistemic authority than women, which contributes to the privileging of masculine voices above feminine ones in society. (McCready and Winterstein, 2017) model this via the assumption that women are generally, on a heuristic and stereotypical basis, assigned a lower degree of reliability than men (as modeled using Bayesian methods in a way summarized in the full paper); equating reliability with epistemic authority, we then expect differences in masculine and feminine speech strategies, which will utilize high degrees, and try to overcome low degrees, of epistemic authority.

What happens when one is deemed to have authority in a social situation? One is able to push one's judgements, and see to their realization, even at the expense of other conflicting possibilities. In conversation, the authoritative speaker can expect that his assertions will be believed, downplay or reject the assertions of others, and to some extent dictate the topic of conversation. If men in general have more epistemic authority than women, it is to be expected that they would use particles which reflect this status: particles like *ge3/laa3* which strengthen assertions, or particles like *ze1* which push aside the assertions of others. Conversely, if women have less epistemic authority, they would be expected to use particles which work to overcome this 'position of weakness,' such as the softener *aa3* or particles such as *me1* that indicate low prior belief, but do not outright reject a statement.

Note that these observations are found in actual speech, i.e. most likely reflect the actual status of feminine speakers more than stereotypes about them. It is to these stereotypes we now turn.

2 The fictional speech of women: directness and *Sajiao*

The results mentioned above are not perfectly in line with those mentioned in the existing literature.

Notably, the particle *zek1*, described by (Chan, 2002) as feminine, does not come out as significantly gender biased in our analysis. We begin by discussing the particle *zek1*. Then we relate it to the larger question of *sajiao*, a recurring concept when discussing Chinese gendered speech.

2.1 *Zek1* as a fiction particle

(Chan, 2002) used transcripts of TV dramas as her main source of data to argue in favor of a feminine bias in the distribution of *zek1*. She and others (Matthews and Yip, 2011) describe *zek1* as a particle which is preferably used in intimate situations, with close friends and family. Most speakers also share the intuition that *zek1* has feminine overtones. A probe of the Mid-20th Century Cantonese Corpus (Chin, 2015), which contains transcription of Cantonese movies produced between 1943 and 1970, gives results comparable to those reported by Chan. There, the frequency of *zek1* is biased toward feminine speakers who use it on average in 1.75% of their utterances, significantly more than male speakers who use it ($\chi^2 = 141.59; df = 1; p < 0.001$).

The corpus data used for our analysis consist of conversations between close friends, family members or close colleagues, i.e. situations similar to the ones depicted in the TV and movie transcripts. The main difference is that our corpus is authentic data, while the transcripts are fiction. We argue that these TV and movie transcripts exhibit stereotypes about women's speech which only partially overlap with genuinely naturalistic data. These stereotypes are founded on some preconceptions about how women and men speak, and fictional speech is very likely to reflect them, especially when aiming at being natural.

There is no reason to consider *zek1* as restricted to intimate situations, as commonly assumed. Its use is most natural in such situations because the speech acts it marks is better directed at intimates. But *zek1* appears in other contexts, for example if a speaker wants to indicate his impatience. Regarding the purported gender bias of *zek1*, we claim that it is related to a more general stereotype about women in Chinese societies, one which is best represented through *sajiao*, which we turn to now.

2.2 Sajiao

The discussion of gendered speech in Chinese languages regularly mentions *sajiao* (撒嬌, Cant. *saat3giu1*, now SJ). SJ is typically described as a communicative style used to wheedle one's audience. The actual definition of SJ is blurry at the edges: it is often defined as an imitation of children's speech (Farris, 1994; Chan, 2002; Yueh, 2013), but at the same time children are described as doing SJ too. However, all descriptions agree on the fact that SJ is stereotypically associated with feminine speakers. SJ is frequently mentioned in popular media: there are books, TV series and movies (typically romantic comedies) about "women who do SJ", and SJ is discussed in talk shows, news outlets etc. (Yueh, 2013) The linguistic correlates of SJ have not been thoroughly investigated; frequently mentioned features include syllable lengthening, hypocoristic reduplication and the use of some SFP.

We will define SJ in the following way. SJ is a communicative strategy aiming at getting some benefit from a potentially unwilling addressee. The speaker presents themselves as being unable to obtain by themselves whatever they want, i.e. they use a display of what (de Certeau, 1984) calls "absence of power". The speaker signals this by using child-like features in a way that is meant to be recognized by the addressee.

The strategy works by lifting certain societal stereotypes to attention. In particular, as with many other societies, men hold societal power in traditional Chinese culture, and are expected to be providers; women, on the other hand, are normatively lacking in power and dependent on men for various things required for life. It should be stressed that these roles are only normative; we can think of them as *ideologies* in the sense of (Stanley, 2015), in that they offer something like an organizing framework for thinking about social relations, but as *flawed ideologies* in that the framework they offer systematically subordinates women. Ideologies are like other sorts of ideas and social frameworks in that they can be deployed strategically; we view *sajiao* as a linguistic device implementing a strategic use of the flawed ideology which assigns women a dependent and subordinate role in traditional Chinese culture (hereafter CGI for 'Chinese Gender Ideology').

Our leading idea is that SJ offers the addressee a chance to assume the role of a provider. In the case of male speakers, this is aligned with the societal stereotype assigned to men, who sterotypically are holders of societal power, heads of household, etc. If the addressee endorses the stereotypes present

in CGI, they will be likely to yield to the speaker's desires after the sterotypes enter attention. In turn, as long as a female speaker is willing to tacitly endorse CGI, it is a good strategic move to make use of SJ, as it gives an increased likelihood of satisfying other demands in the right situation.

Concretely, we can model this view in a game-theoretic setting as follows. Suppose that a female speaker a wants to make a request of a male speaker b and that satisfying the request will yield a payoff of α for a and a utility of β for b, which will be lower than the utility assigned when the request is refused, due to the effort expended. We model the effect of SJ by taking (Quinley, 2012) as a starting point. There, requests are modeled as trust games, in which one player makes a request and the other chooses whether not to honor it based in part on the form of the request: a more polite request gives the speaker additional positive face, in the sense of (Brown and Levinson, 1987), and a sufficiently polite request (over a sufficient number of game iterations, in a repeated game setting) gives incentive to satisfy the request in order to gain face. Thus, SJ can be viewed as a special kind of positive politeness; by emphasizing her subordinate status, a (female) speaker can give positive face to her (male) interlocutor, though she simultaneously lowers her status. This can be modeled by assigning a utility penalty of γ to a, while a utility benefit of ρ accrues to b: in the SJ case, then a's payoff is $\alpha - \gamma$, while b receives $\beta + \rho$.

The following extensive form game shows a sample case. The game models a request by a for b to do something. There are two cases: one where a does not use *sajiao*, and one where she does. If a does not, then, if b accepts her request, she receives a payoff of 4; if her request is rejected, she gets only 2. Conversely, b gets 4 if he ignores the request and 2.5 otherwise (a small utility bonus is given for getting face by being asked to do something and being able to do it, but penalty exacted for the cost of the action). If a chooses to use SJ, she takes a face penalty of 1 for utilizing a flawed ideology which subordinates her; but b's incentive to fulfill her request is larger than before, increasing the rationale for b to choose to do what she asks. In this sense, the use of SJ is a rational, strategic move, with two caveats: first, if b doesn't subscribe to CGI, or even opposes it, negative consequences ensue for a (notably that her request will likely be rejected even though she's subordinated herself), and, second, that it's still rational for b to reject the request in a one-off setting, meaning that repeated games are required (as already pointed out by (Quinley, 2012) for the non-SJ case; see also (McCready, 2015) for the general issue of cooperation).

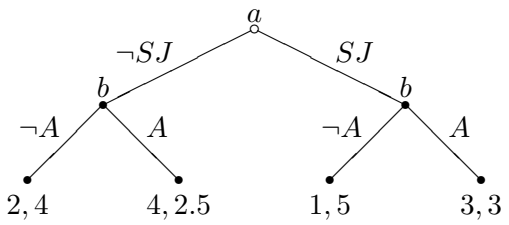

Figure 1: Sajiao speech

This definition is rather restrictive and excludes cases which have been analyzed as SJ, e.g. those discussed by (Yueh, 2013). It does not characterize SJ by its form, meaning that the use of specific intonation or SFP is not sufficient to qualify a communicative act as SJ. This definition makes SJ compatible with both adult and children speakers, and is not gender specific. That being said, in order to use SJ, one's status should be compatible with presenting oneself as powerless and the status of the addressee should be compatible with being presented as a "provider". Thus, it is the existence of CGI that makes SJ effective, which in turn accounts for why SJ is mostly used by women and children and targeted at men or parents.

3 Conclusion

In this paper, we looked at two registers related to feminine speech. The SJ register was analyzed as a strategic use of cultural ideologies to maximize one's benefits in repeated interactions, similar to politeness, but dependent on an acceptance of these ideologies. The other register we analyzed through the distribution of SFPs. This one contrasts with SJ in that it does not exploit ideologies for strategic benefit, but rather shows how women use linguistic tools to deal with the personae assigned to them in traditional society, i.e. their assigned low degree of epistemic authority. We also argued that SJ is more widely present in fiction and in people's representation of feminine speech than in natural contexts. This naturally follows if one sees SJ as a conscious communicative strategy which speakers, and thus producers of fiction, are sensitive to.

References

Penelope Brown and Stephen C. Levinson. 1987. *Politeness*. Cambridge University Press, Cambridge.

Heather Burnett. Accepted. Signalling Games, Sociolinguistic Variation and the Construction of Style. *Linguistics and Philosophy*.

Marjorie K.M. Chan. 1998. Gender differences in the Chinese language: A preliminary report. In Hua Lin, editor, *Proceedings of the Ninth North American Conference on Chinese Linguistics*, volume 2, pages 35–52, Los Angeles. University of Southern California, GSIL Publications.

Marjorie K.M. Chan. 2000. Gender, Society and the Chinese Language. In Baozhang He and Wenze Hu, editors, *Proceedings of the Eleventh North American Conference on Chinese Linguistics (NACCL-11)*, Cambridge, MA. Harvard University.

Marjorie K.M. Chan. 2002. Chinese: Gender-related use of sentence-final particles in Cantonese. In Marlis Hellinger and Hadumod Bussmann, editors, *Gender Across Languages: The Linguistic Representation of Women and Men.*, volume 2, pages 57–72. John Benjamins, Amsterdam.

Andy Chin. 2015. A linguistics corpus of mid-20th century hong kong cantonese. Department of Linguistics and Modern Language Studies, The Hong Kong Institute of Education, Retrieved 23/3/2015.

M. de Certeau. 1984. *The practice of everyday life*. University of California Press, Berkeley.

Mary S. Erbaugh. 1985. Sentence final particles as an Asian areal feature. In Scott De Lancey and Russell S. Tomlin, editors, *Proceedings of the First Annual Meeting of the Pacific Linguistics Conference*, pages 84–96, Eugene, Oregon. University of Oregon.

Catherine S. Farris. 1994. A Semiotic Analysis of Sajiao As a Gender Marked Communication Style in Chinese. In Marshall Johnson and Fred Y. L. Chiu, editors, *Unbound Taiwan: Close-ups from a distance.*, volume 8, pages 1–29. Center for East Asian Studies, The University of Chicago.

Miranda Fricker. 2009. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

Yurie Hara and Eric McCready. in press. Particles of (un)expectedness: Cantonese wo and lo. In *Selected post-proceedings of LENLS 2015*. Springer.

Helen Kwok. 1984. *Sentence Particles in Cantonese*. Center of Asian Studies, University of Hong Kong.

Timothy Light. 1982. On being 'de-ing'. *Computational Analyses of Asian and African Languages*, 19:21–49.

Kang Kwong Luke and May L.Y. Wong. 2015. The Hong Kong Cantonese Corpus: Design and Uses. *Journal of Chinese Linguistics*, 25:312–333.

Kang Kwong Luke. 1990. *Utterance Particles in Cantonese Conversation*. John Benjamins Publishing Company, Amsterdam, Philadelphia.

Stephen Matthews and Virginia Yip. 2011. *Cantonese, A Comprehensive Grammar*. Routledge, 2nd edition.

Eric McCready and Grégoire Winterstein. 2017. Negotiating epistemic authority. In Kurahashi S. et al, editor, *New Frontiers in Artificial Intelligence*. Springer, Berlin. to be published.

Eric McCready. 2015. *Reliability in Pragmatics*. Oxford University Press, Oxford.

Jason Quinley. 2012. Trust games as a model for requests. *New Directions in Logic, Language and Computation*, pages 221–233.

Jason Stanley. 2015. *How Propaganda Works*. Princeton University Press, Princeton.

Hsin-I Sydney Yueh. 2013. Body Performance in Gendered Language. Deconstructing the Mandarin Term Sajiao in the Cultural Context of Taiwan. *Journal of Theories and Research in Education*, 8(1):159–182.

Achim Zeileis, David Meyer, and Kurt Hornik. 2007. Residual-based shadings for visualizing (conditional) independence. *Journal of Computational and Graphical Statistics*, 16(3):507–525.