

## **UC Merced**

# **Proceedings of the Annual Meeting of the Cognitive Science Society**

### **Title**

Biological Males' and 'Trans(gender) Women': Social Considerations in the Production of Referring Expressions

### **Permalink**

<https://escholarship.org/uc/item/8b5481wb>

### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

### **Authors**

Papineau, Brandon

Degen, Judith

### **Publication Date**

2024

Peer reviewed

# ‘Biological Males’ and ‘Trans(gender) Women’: Social Considerations in the Production of Referring Expressions

Brandon Papineau, Judith Degen

{branpap, jdegen}@stanford.edu

Department of Linguistics, 460 Jane Stanford Way  
Stanford, CA, 94305

## Abstract

Understanding referring expression generation has long been of interest to psycholinguistics, pragmatics, and sociolinguistics. Experimental data in the former two has shown that referring expression generation is modulated by both pragmatic and cognitive considerations, and the latter suggests that referring expressions have social meaning beyond their literal referential utility. This project integrates these three accounts by extending Burnett (2017)’s socially-enriched implementation of the Rational Speech Act (RSA) framework to account for variation in referring expressions used to denote transgender women in two politically opposed media corpora. Our findings highlight the utility of the RSA framework in explaining socially-modulated variation while also accounting for pragmatic and cognitive considerations. Finally, this paper contributes to growing literatures that address the relationship between (alt-)right ideologies about gender and language by highlighting the use of bioessentialist language such as *biological male* in the propagation of anti-trans rhetoric in the United States.

**Keywords:** language and gender; language production; language and politics; media linguistics; Rational Speech Act

## Introduction

In 2023, *Mother Jones* reported on the leaking of over 2600 pages of emails between South Dakota Republican Fred Deutsch and legal, political, and media associates which dealt with an attempt to pass anti-transgender legislation in the state. Contained in these emails is the following:

“I would not want to see anyone on our side intentionally put those terms [cisgender, transgender, non-binary] into the record. Using them surrenders the language. Language frames the debate. If the other side’s language frames the debate, we lose.”

This quote exemplifies the crucial role of language in the political sphere, and the ways in which it can structure discourse and be exploited to either uphold and reproduce or challenge hierarchical power structures. In this case, language that ratifies the social gender of trans and non-binary folks is explicitly avoided so as to not legitimize the distinction between sex assigned at birth and social gender. This accords with recent investigations into the relationship between right-wing politics and gender which have shown that traditional gender ideologies are strongly correlated with (alt-)right political leanings (Christley, 2022; Corredor, 2019; Gökariksel et al., 2019), and that the proliferation of gender-neutral language has emerged as a point of contention for

those on the political right (Erdocia, 2022; Papineau et al., 2022).

One crucial linguistic device used to uphold the aforementioned power and social relations is the exploitation of referring expressions, or expressions used to denote a particular entity in a given discourse. Findings in sociolinguistics highlight this, with expressions such as *dude* (Kiesling, 2004), *tongzhi* (a Chinese term used to denote members of the queer community which originates from the Chinese word ‘comrade’; Wong, 2005), and *Arabian stallion* (Meier-Vieracker, 2023) having been described as being imbued with indexical power, or ideological links between linguistic forms and the social meanings attached to those terms (Eckert, 2008). These forms and their attached social indexicalities then bundle together with other form-meaning pairs to create *personae*, or interaction-based identities. As such, these terms are employed not only to construct the personae of those employing them, but also to denigrate and delegitimize others. This project contributes to this line of inquiry by examining how the terms *biological male*, *transgender woman*, and *trans woman*, used to denote transgender women, vary within and between two politically-opposed news corpora (the alt-right *Breitbart* and the queer publication *PinkNews*). We argue that these expressions, like others, are used to simultaneously construct authors’ personae and (de)legitimize trans identities.

Outwith sociolinguistics, referring expression generation has also received attention in both the pragmatics and psycholinguistic literatures. The former of these has been primarily concerned with the notion of informativity, or how likely a listener is to recover an intended meaning given an utterance, and has found for example that speakers sometimes redundantly modify their utterances (Dale & Viethen, 2009; Degen et al., 2020; Gatt et al., 2013), and do so in systematic ways that belie cognitive underpinnings such as the preference to include atypical information or redundant color information. (Kursat et al., 2022; Tarenskeen et al., 2015; Van Gompel et al., 2019; Westerbeek et al., 2015).

Psycholinguistics, on the other hand, has focused on the effect of utterance cost on referring expression generation, identifying cognitive costs such as utterance length (Degen et al., 2013; Karimi et al., 2014) and the relative accessibility of terms (Fukumura & Van Gompel, 2011; Vogels, 2014) as factors in generation.

Under a sociolinguistic account, we expect that *trans*

woman and transgender woman should occur frequently in the *PinkNews* corpus, while only *biological male* should occur in *Breitbart*, as the terms serve specific social goals by constructing identities that accept or reject the distinction between sex assigned at birth and social gender, respectively. A pragmatic account would predict the distribution of terms to be determined by their referential utility, which appears to be the same between each of the three expressions, as they all serve to pick out trans women. Crucially, ‘biological male’ is not used in the data to refer to cisgender men, suggesting its in-context functional equivalence with the other two terms. Finally, a psycholinguistic account would predict that the least costly term will be favored, in this case *trans woman* as operationalized by COCA frequency counts (Davies, 2008-, see below).

To account for these varying factors and explain the observed variation in our corpora, we draw on the Rational Speech Act framework (Frank & Goodman, 2012; Goodman & Frank, 2016), and in particular Burnett (2017)’s model of identity construction. We show that only an integrated account of referring expression generation that factors in sociolinguistic, pragmatic, and psycholinguistic considerations can accurately account for the variation we see within and between the two publications. In doing so, this project also contributes to a burgeoning literature that seeks to identify the relationship between political identity and language variation (Burnett & Pozniak, 2021; Hall-Lew & Trousdale, 2020; Hall-Lew et al., 2010; Krivoruchko, 2008; Papineau et al., 2022), as well as one that seeks to identify individual differences in referring expression generation (Baltaretu & Ferreira, 2016; Hendriks, 2016; Viethen & Dale, 2010).

## The Study

### The Corpora

The data in this project comes from two purpose-built corpora, constructed to represent the linguistic practices of both the alt-right and queer communities as they regard trans and non-binary identities<sup>1</sup>. We chose a queer-focused corpus for comparison to avoid using the political ‘middle’ as the standard. We instead consider queer voices as the authoritative standard to which practices of the political right should be contrasted, as queer authors are more likely to ratify transgender identities and thus serve as the right’s mirror image in this regard.

The political (alt-)right is represented by *Breitbart*, known for being one of the bastions of the American alt-right, described by other researchers as a ‘central node on key issues within an informational network of right-wing ideological views’ (Heft et al., 2020, p. 24). The site’s news is shared disproportionately frequently on social media networking sites such as Facebook (Economist, 2020), and claims to be among the most engaged with Instagram and Facebook pages (Breitbart, 2019). Previous research into the narratives constructed

<sup>1</sup>All materials and data are available at <https://github.com/BranP/ap/Transgender-Referring-Expressions>

at *Breitbart* have placed it as something of an ideological middle ground between extreme right and centre-right media, as it presents ‘a more radical version of *Fox News* that tries to solidify a common identity among fragmented far-right groups ... but simultaneously tries to connect with the mainstream right and potentially influence its agenda via shared topics’ (Kaiser et al., 2020, p. 437).

The queer community in this study is represented by *PinkNews*, a queer publication originally based out of the UK. *PinkNews* covers a wide range of topics, including human interest stories highlighting the achievements of queer people, as well as stories related to discrimination, violence, and oppression of queer peoples around the world.

Python scripts were used to extract the headlines and text of all articles tagged ‘transgender’ on the two sites between their relative inceptions and November 11th, 2022 (inclusive). Summary statistics of this process are presented in Table 1.

	Breitbart	PinkNews
Founded	2007	2005
Articles Pulled	2,936	1,984
Words in Corpus	8,831,627	5,829,504

Table 1: Statistics for the two corpora

### Expressions of Interest

***Breitbart*** *Breitbart* employs language that highlights the sex assigned at birth of trans people, as in *biological male*. In the following example about NCAA trans swimmer Lia Thomas and professional disc golfer Natalie Ryan, the term is used to further the narrative that trans women are just ‘men in dresses’ who seek to dethrone (cisgender) women from various accolades. This implicitly endorses the idea that gender is biologically ‘simple’ and inextricable from sex assigned at birth (a *bioessentialist* belief associated with anti-trans rhetoric; Martin and Rahilly, 2023), as evinced by the use of the verb ‘claiming’ which implicates a failure of social gender transition. These narratives are repeated throughout *Breitbart* in order to bedevil and delegitimize trans identities.

‘In shades of biological male Lia Thomas’s domination of women’s college swimming last season, the world of professional women’s disc golf now finds itself beset with a biological male claiming to be a transgender woman who is knocking natural-born women out of top money all across the country’. (Huston, 2022)<sup>2</sup>

***PinkNews*** Turning now to *PinkNews* with a quote discussing Thomas’ nomination for the 2022 GLAAD awards, we see that the referring expressions used to identify her differ greatly from those in *Breitbart*. Thomas is positioned as

<sup>2</sup><https://www.breitbart.com/sports/2022/09/29/biological-male-dominating-opponents-womens-pro-disc-golf/>

an advocate and champion of transgender rights, using her unfortunate position as the object of controversy to advocate on behalf of all ‘trans people’. Importantly, she is not identified as ‘biological male Lia Thomas’, but rather as ‘swimmer Lia Thomas’, foregrounding her achievements rather than her sex assigned at birth. This ratification of her gender and avoidance of sex-based language is supported by the use of the gender-congruent pronouns (‘her’, ‘herself’), as well as in the crucial referring expression ‘trans women’. Finally, the use of language such as ‘allowance’ and ‘excluded’ positions Thomas and other trans women not as the perpetrators of inequality in the sporting world, but rather as the victims of it—the opposite framing to the one presented in *Breitbart*.

Swimmer Lia Thomas has used her platform to advocate for the allowance of trans people in sporting events, having been excluded herself after an effective ban on trans women in swimming competitions. (Hansford, 2022)<sup>3</sup>

## Results

Monroe et al. (2008)’s *Fighting Words Algorithm* (FWA) was used to identify bigrams most closely associated with each of the two corpora; the results of this analysis are provided in Table 2. Crucially, the referring expression *biological males* appears as one of the bigrams most closely associated with *Breitbart*, while *PinkNews* was associated with the strings *transgender woman*, *trans woman*, and *trans women*. This finding suggests that referring expressions used to denote transgender women are socially or politically motivated, a suggestion corroborated by the descriptive patterns of referring expression distribution, provided in Fig. 1.

*Biological male* appears almost exclusively in the *Breitbart* data, while *PinkNews* employs language that ratifies the gender identities of the individuals they discuss (i.e. *trans(gender) woman*), while simultaneously avoiding sex-based language. Not predicted by a sociolinguistic account, however, is the fact that *transgender woman* is more common in the *Breitbart* data than *biological male*, and *trans woman* trails *biological male* by only a little. We identified the source of this unpredicted result as the heavy use of quotations and embedded media in *Breitbart*. A random sample of 50 tokens of each term that referred to transgender women in *Breitbart* was coded for whether or not they occurred in quotations or embedded media. These proportions were used to estimate the genuine WPM values of referring expressions in *Breitbart*, the results of which are provided in Table 3.

The result of this process is a qualitative pattern in *Breitbart* such that *biological male* is the most genuine to the publication, followed closely by *transgender woman* and then *trans woman* at a distance. All 8 occurrences of *biological male* in *PinkNews* occurred in quotations, and so were removed from analysis.




Figure 1: Distribution of referring expressions denoting trans women across the two corpora

## Model of Referring Expressions

Having now provided a descriptive account of the variation we see in the corpora, we turn to modeling this variation using the Rational Speech Act (RSA) framework. Crucially, this framework allows us to incorporate language users’ beliefs about the state of the world, such as the likely messages or social beliefs an interlocutor is likely to be conveying in a particular interaction. This in turn allows us to test whether or not social considerations assist in modeling the observed variation.

## Social Signaling Games

The RSA framework (Degen, 2023; Frank & Goodman, 2012; Goodman & Frank, 2016) provides a way of understanding interactions between rational agents. At its core, RSA posits that a *pragmatic speaker* ( $S_1$ ) reasons about the *literal listener* ( $L_0$ ), who serves to capture the interpretation of possible forms based on their literal semantics and disallow untrue utterances (1). For example, if the utterance *lightbulb* is used to denote a transgender woman, the  $L_0$  returns 0 and the pair is not considered as a viable utterance-message pair.

### (1) The Literal Listener

$$P_{L_0}(m|u) = \delta_{m \in [u]} \cdot P(m)$$

The *pragmatic speaker* then reasons about the literal listener and selects an utterance by applying a softmax function to the utterance  $u$ ’s *utterance utility* which balances utterance informativity and the cost associated with its production, as in 2. Under this formulation, the utility ( $U$ ) of an utterance-message pair ( $u, m$ ) is equal to the log probability of a literal listener ( $L_0$ ) recovering the intended meaning ( $m$ ) from a

<sup>3</sup><https://www.thepinknews.com/2022/10/12/pinknews-awards-2022-nominations-united-airlines-lloyds-banking-group/>

Breitbart				PinkNews			
Bigram	n	WPM	Z-Score	Bigram	n	WPM	Z-Score
transgender ideology	687	77.79	13.720	trans people	1834	314.61	-20.481
biological sex	639	72.35	11.728	transgender people	1549	265.72	-13.272
opposite sex	511	57.86	11.181	gender recognition	423	72.56	-10.659
Lia Thomas	319	36.12	9.567	<b>transgender woman</b>	509	87.31	-8.190
puberty blockers	375	42.46	8.815	was shot	193	33.11	-8.150
gender ideology	306	34.65	8.739	<b>trans woman</b>	385	66.04	-8.113
title ix	332	37.59	8.178	recognition act	162	27.79	-7.571
transgender activists	328	37.14	8.118	was found	216	37.05	-6.958
gender identity	1811	205.06	7.849	trans rights	229	39.28	-6.624
<b>biological males</b>	232	26.27	7.555	<b>trans women</b>	499	85.60	-6.580

Table 2: Bigrams most associated with the two corpora. n refers to the raw number of occurrences. WPM refers to the number of occurrences per million words. Z-Scores refer to the output of the Fightin’ Words Algorithm; a more negative score indicates a higher association with the *PinkNews* corpus, while a more positive number indicates an association with *Breitbart*.

Term	Genuine WPM Estimate
<i>biological male</i>	30.98
<i>transgender woman</i>	29.29
<i>trans woman</i>	5.36

Table 3: Genuine WPM estimates for the terms in *Breitbart*

given utterance ( $u$ ) minus the cost of said utterance. The term *transgender woman*, for example, might strike the proper balance between utterance informativity and cost, and thus be presented as a good option for the *pragmatic speaker*.

### (2) The Pragmatic Speaker

$$P_{S_1}(u|m) \propto \exp(\alpha \cdot \ln P_{L_0}(m|u) - \text{cost}(u))$$

The *pragmatic listener* ( $L_1$ ) then exploits their prior beliefs about the world regardless of message (for example, the belief that a *Breitbart* author is unlikely to have a progressive persona), captured as *priors* ( $P(m)$ ), to inform their recovery of intended meaning, along with considerations of what and how the *pragmatic speaker* is likely to be communicating.

### (3) The Pragmatic Listener

$$P_{L_1}(m|u) \propto P_{S_1}(u|m) \cdot P(m)$$

This model has found ample empirical support across a variety of domains, being used to account for phenomena such as free choice disjunction (Champollion et al., 2019), referring expression generation (Degen et al., 2020; Kursat et al., 2022; Waldon & Degen, 2021), and scalar implicature (Degen, 2023; Zhang et al., 2023).

Burnett (2017) extends this framework to variationist sociolinguistics. In this adaptation, the utterance-message pair consists of a social persona (meaning) and the linguistic form used to communicate and recover that persona, rather than a linguistic form and its literal semantic meaning. She exemplifies the framework’s utility in explaining variation in both (ING) variation (e.g. *slaying* vs. *slayin’*) in a non-binary community (Gratton, 2016), as well as the relationship between

gender stereotypes and the perception of male pitch variation (Levon, 2014). We adopt this model, and lay out our assumptions next.

## Model Assumptions

**Referential Utility** We begin by assigning each of the messages (in this case, referring expressions) to the *personae* we take them to be capable of indexing. We assume two parts to the *personae*; in these interactions, referring expressions can index either a conservative or progressive political stance, as well as either a bioessentialist or gender-progressive ideology. Crossing the two social facets results in four possible *personae* (Table 4). Under this formulation, *biological male* can only index a conservative, bioessentialist persona, while *trans woman* can only index its opposite. Because there is no reason to believe that conservatism and bioessentialism are necessarily intertwined, the other two *personae* are possible as well. However, because of the strong indexical values of *biological male* and *trans woman*, these *personae* can only be communicated through the neutral ‘transgender woman’, which is relatively unmarked for these two attributes.

Utterance	Message (personae)
<i>biological male</i>	{conservative, bioessentialist}
<i>transgender woman</i>	{conservative, bioessentialist}, {conservative, gender-prog.}, {progressive, bioessentialist}, {progressive, gender-prog.}
<i>trans woman</i>	{progressive, gender-progressive}

Table 4: Possible *personae* communicated by referring expressions

**Priors** We assume three separate sets of priors, one each for an agnostic listener (i.e., a listener with no preconceived notions about *Breitbart* or *PinkNews* and their leanings), a

*Breitbart* prior set that captures the expectations of *Breitbart* readers (BB), and the same for readers of *PinkNews* (PN). While a naive reader will not expect any particular persona based on the publication, we assume that *Breitbart* readers will generally expect author personae that are conservative and bioessentialist and that *PinkNews* readers will generally expect progressive, gender-progressive ones; *PinkNews* espouses left-wing leanings and has explicitly campaigned against Conservative politicians in the UK and GOP politicians in the US. Additional prior mass is assigned to those personae which are politically but not gender-ideologically congruent with publication leaning (Table 5).

Persona	Naive	BB	PN
{conservative, bioessentialist}	0.25	0.9	0.001
{conservative, gender-prog.}	0.25	0.098	0.001
{progressive, bioessentialist}	0.25	0.001	0.098
{progressive, gender-prog.}	0.25	0.001	0.9

Table 5: Priors for each of the assumed reading audiences

**Cost** We assume that the less frequent a term is, the more costly / less accessible it is to a language user. Cost considerations based on COCA frequency values are presented in Table 6 (Davies, 2008-).

	Log Frequency
<i>biological male</i>	3
<i>transgender woman</i>	4.8
<i>trans woman</i>	5.4

Table 6: Values used to assign cost values in the model

Under this cost function *trans woman* is the least costly form, but note that the differences between costs for the three terms are relatively small, and that it is unlikely that *Breitbart* and *PinkNews* actually have similar frequency exposure to the the three terms (e.g., *Breitbart* readers likely encounter *biological male* more frequently than do *PinkNews* readers). Future work is will explore the full possible cost space, including accounting for differences in linguistic exposure and experience.

**Alpha** Finally, we assume an  $\alpha$  parameter of 1.

## Model Results

Implementing these assumptions into Burnett (2017)’s model results in the predictions presented in Fig. 2 and Table 7, the latter of which contains the observed proportions of utterances alongside the predicted proportions. The numbers indicated on the x-axis of Fig. 2 represent the probability that a pragmatic speaker will employ a specific message given an intended persona. For example, a pragmatic speaker in a *Breitbart* context is expected to use the term *biological male* 0.52% of the time to communicate a conservative, bioessentialist persona.

We can essentially reject the naive priors out of hand, as they predict that a writer indexing a conservative, bioessentialist persona will use *biological male* 80% of the time, with the remaining 20% of occurrences being *transgender woman*, which is not a good prediction of the *Breitbart* data– nor is its inverse a good predictor of the *PinkNews* data.

The publication-specific models, however, are good predictors of the observed proportions in the respective corpora (Table 7). In the case of *Breitbart*, the model predicts the correct qualitative pattern, but fails to capture the mere presence of *trans woman* in the corpus. Post-hoc examination of the data indicates that many of these realizations occur in constructed dialogue (Tannen, 1986), or dialogues and sentiments attributed to other social actors, and may be additionally classed as not genuine to *Breitbart*. We leave the exploration of the role of constructed dialogue in this discourse for future work.

The *PinkNews* predicted values are nearly perfect, on the other hand, and require no such additional explanation. The result is an extremely satisfying one: merely mirroring the ideological priors between two ideologically-opposed sources of media is able to account for the variation of referring expressions within and between them. Crucially, this is only achieved upon integrating social beliefs into the pragmatic reasoning game, as purely pragmatic, psycholinguistic, or sociolinguistic accounts fail to make accurate predictions.

The patterns above make it clear that an integrated approach to explaining socially-conditioned variation is paramount. While the indexicality account explains the broad pattern that *PinkNews* prefers *trans woman* and *Breitbart* prefers *biological male*, it alone is unable to account for the presence of *transgender woman* appearing in *Breitbart*, as well as the high occurrence of that term in both corpora. While a purely referential understanding of pragmatic reasoning in referring expression generation similarly fails to predict the patterns, employing this framework as a way of modeling indexicality probabilistically allows us to better capture the observed proportions of referring expressions when contextually-driven priors are incorporated. Finally, a cost-based understanding of referring expression generation alone predicts much higher use of *trans woman* than is observed, but can be successfully integrated into the fuller RSA model. Only with the integration of all three frameworks can we best account for the variation in referring expressions in *Breitbart* and *PinkNews*.

## Discussion

While the referring expression generation literature is a rich one that has previously considered information redundancy (Degen et al., 2020), property atypicality (Westerbeek et al., 2015), relative surprisal (Stegemann-Philipps et al., 2021), and cognitive salience (Dale & Viethen, 2009; Gatt et al., 2013), among other considerations, work on individual or social variation is relatively infrequently incorporated into these models (though cf. Baltaretu and Ferreira (2016), Hendriks




Figure 2: Predicted production probabilities for the three terms (color) to communicate each of the 4 personae (x-axis), under three different sets of priors (facet).

Utterance	Breitbart (Observed)	Breitbart (Predicted)	PinkNews (Observed)	PinkNews (Predicted)
<i>biological male</i>	0.47	0.52	0	0
<i>transgender woman</i>	0.45	0.48	0.46	0.47
<i>trans woman</i>	0.08	0	0.54	0.53

Table 7: Observed and model-predicted proportions of referring expressions in the two corpora. Predictions made by the models are for the persona we assume is being communicated by each publication (i.e. *conservative, bioessentialist* for *Breitbart* and *progressive, gender-progressive* for *PinkNews*).

(2016), and Viethen and Dale (2010)).

By extending Burnett (2017)’s framework of signaling games for social meaning to the lexical realm, we have shown that the RSA model lends us the tools for incorporating ideological and social goals and identities into this line of investigation. Doing so allows us to best account for the variation in natural data observed in our socio-politically polarized news sites, and highlights the need for integrated accounts of language production that include pragmatic, cognitive, and social considerations.

In doing so, we also contribute to explaining politically-driven language variation, which has received relatively little attention from sociolinguists (though cf. Burnett and Poznaniak, 2021; Hall-Lew et al., 2010; Krivoruchko, 2008; Papineau et al., 2022). This literature suggests that political ideology and identity do serve as predictors of linguistic behavior, and we argue that it would serve sociolinguistics well to consider these dimensions along with traditional social considerations.

We plan to investigate the perception side of this puzzle by examining how well individuals recover the assumed in-

dexicalities we have posited herein; such data will allow for the development of a fuller model incorporating the Pragmatic Listener. We also acknowledge that the use of only politically-driven news corpora leaves unanswered how those in the ‘mainstream media’ behave, and future work will incorporate more politically moderate media.

Finally, this paper contributes to a literature which has examined the role of the political (alt-)right in promotion of traditional gender values which threaten the rights of queer, trans, and non-binary people, as well as those looking for access to abortion and equal rights (Christley, 2022; Corredor, 2019; Erdocia, 2022). This paper serves to show that a holistic approach that draws on a range of disciplines will only enrich our understanding of the relationships between right-wing ideologies, gender, and language. We hope that this paper will encourage others to engage critically with the language of the political right, and to take social and ideological variation into account in studies of denigration, ratification, and referring expression generation.

## Acknowledgments

We would like to thank Katerine Hilton, Beth Levin, Nandi Sims, and Rob Podesva, as well as the members of the ALPS and Interactional Sociophonetics Labs at Stanford for helpful discussion. We also thank the four anonymous CogSci reviewers whose feedback has shaped this paper. We would also like to extend our thanks to Adolfo Hermsillo, Alexia Hernandez, Anthony Velasquez, Bonnie Krejci, Brandon Waldon, Evelyn Fernández-Lizárraga, Lewis Esposito, Madelaine O'Reilly-Brown. While all of these folks have contributed to this paper in some way, shape, or form, any remaining mistakes or shortcomings are our own.

## References

- Baltaretu, A., & Ferreira, T. C. (2016). Task demands and individual variation in referring expressions. *Proceedings of the 9th International Natural Language Generation conference*, 89–93.
- Breitbart. (2019). Who Breitbart is...really. <https://media.breitbart.com/media/2019/11/about-breitbart-news.pdf>
- Burnett, H. (2017). Sociolinguistic interaction and identity construction: The view from game-theoretic pragmatics. *Journal of Sociolinguistics*, 21(2), 238–271.
- Burnett, H., & Pozniak, C. (2021). Political dimensions of gender inclusive writing in Parisian universities. *Journal of Sociolinguistics*, 25(5), 808–831.
- Champollion, L., Alsop, A., & Grosu, I. (2019). Free choice disjunction as a rational speech act. *Semantics and Linguistic Theory*, 29, 238–257.
- Christley, O. R. (2022). Traditional gender attitudes, nativism, and support for the radical right. *Politics & Gender*, 18(4), 1141–1167.
- Corredor, E. S. (2019). Unpacking “gender ideology” and the global right’s antigender countermovement. *Signs: Journal of Women in Culture and Society*, 44(3), 613–638.
- Dale, R., & Viethen, J. (2009). Referring expression generation through attribute-based heuristics. *Proceedings of the 12th European Workshop On Natural Language Generation (ENLG 2009)*, 58–65.
- Davies, M. (2008-). The corpus of contemporary american english (coca). <https://www.english-corpora.org/coca/>
- Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, 9, 519–540.
- Degen, J., Franke, M., & Jager, G. (2013). Cost-based pragmatic inference about referential expressions. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 35(35).
- Degen, J., Hawkins, R. D., Graf, C., Kreiss, E., & Goodman, N. D. (2020). When redundancy is useful: A Bayesian approach to “overinformative” referring expressions. *Psychological Review*, 127(4), 591.
- Eckert, P. (2008). Variation and the indexical field. *Journal of Sociolinguistics*, 12(4), 453–476.
- Economist, T. (2020). Facebook offers a distorted view of American news. <https://www.economist.com/graphic-detail/2020/09/10/facebook-offers-a-distorted-view-of-american-news>
- Erdocia, I. (2022). Language and culture wars: The far right’s struggle against gender-neutral language. *Journal of Language and Politics*, 21(6), 847–866.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Fukumura, K., & Van Gompel, R. P. (2011). The effect of animacy on the choice of referring expression. *Language and Cognitive Processes*, 26(10), 1472–1504.
- Gatt, A., Krahmer, E., Van Gompel, R., & Van Deemter, K. (2013). Production of referring expressions: Preference trumps discrimination. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 35(35).
- Gökarıksel, B., Neubert, C., & Smith, S. (2019). Demographic fever dreams: Fragile masculinity and population politics in the rise of the global right. *Signs: Journal of Women in Culture and Society*, 44(3), 561–587.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818–829.
- Gratton, C. (2016). Resisting the gender binary: The use of (ing) in the construction of non-binary transgender identities. *University of Pennsylvania Working Papers in Linguistics*, 22(2), 7.
- Hall-Lew, L., Coppock, E., & Starr, R. L. (2010). Indexing political persuasion: Variation in the Iraq vowels. *American Speech*, 85(1), 91–102.
- Hall-Lew, L., & Trousdale, G. (2020). Breksit or Bregzit: When political ideology drives language ideology. *University of Pennsylvania Working Papers in Linguistics*, 26(2), 11.
- Heft, A., Mayerhöffer, E., Reinhardt, S., & Knüpfer, C. (2020). Beyond Breitbart: Comparing right-wing digital news infrastructures in six Western democracies. *Policy & Internet*, 12(1), 20–45. <https://doi.org/https://doi.org/10.1002/poi3.219>
- Hendriks, P. (2016). Cognitive modeling of individual variation in reference production and comprehension. *Frontiers in Psychology*, 7, 506.
- Kaiser, J., Rauchfleisch, A., & Bourassa, N. (2020). Connecting the (far-) right dots: A topic modeling and hyperlink analysis of (far-) right media coverage during the US elections 2016. *Digital Journalism*, 8(3), 422–441.
- Karimi, H., Fukumura, K., Ferreira, F., & Pickering, M. J. (2014). The effect of noun phrase length on the form of referring expressions. *Memory & Cognition*, 42, 993–1009.



- Kiesling, S. F. (2004). Dude. *American Speech*, 79(3), 281–305.
- Krivoruchko, J. G. (2008). Prepositional wars. *General Editor*, 191.
- Kursat, L., Waldon, B., Ergin, R., & Degen, J. (2022). Evaluating models of referring expression production on an emerging sign language. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 44(44).
- Levon, E. (2014). Categories, stereotypes, and the linguistic perception of sexuality. *Language in Society*, 43(5), 539–566.
- Martin, K., & Rahilly, E. (2023). Value frames in discourse supporting transgender athlete bans. *Discourse & Society*, 34(6), 732–751.
- Meier-Vieracker, S. (2023). Racist discourse in a German far-right blog: A corpus-driven approach using word embeddings. *Discourse & Society*, 09579265231204510.
- Monroe, B. L., Colaresi, M. P., & Quinn, K. M. (2008). Fightin' words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4), 372–403.
- Papineau, B., Podesva, R., & Degen, J. (2022). 'Sally the congressperson': The role of individual ideology on the processing and production of English gender-neutral role nouns. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 44(44).
- Stegemann-Philipps, C., Butz, M. V., Winkler, S., & Achimova, A. (2021). Speakers use more informative referring expressions to describe surprising events. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 43(43).
- Tannen, D. (1986). Introducing constructed dialogue in Greek and American conversational and literary narrative. *Direct and Indirect Speech*, 31, 311–332.
- Tarenskeen, S., Broersma, M., & Geurts, B. (2015). Over-specification of color, pattern, and size: Salience, absoluteness, and consistency. *Frontiers in Psychology*, 6, 1703.
- Van Gompel, R. P., Van Deemter, K., Gatt, A., Snoeren, R., & Krahmer, E. J. (2019). Conceptualization in reference production: Probabilistic modeling and experimental testing. *Psychological Review*, 126(3), 345.
- Viethen, J., & Dale, R. (2010). Speaker-dependent variation in content selection for referring expression generation. *Proceedings of the Australasian Language Technology Association Workshop 2010*, 81–89. <https://aclanthology.org/U10-1013>
- Vogels, J. (2014). *Referential choices in language production: The role of accessibility* [Doctoral dissertation, Tilburg University].
- Waldon, B., & Degen, J. (2021). Modeling cross-linguistic production of referring expressions. *Proceedings of the Society for Computation in Linguistics 2021*, 206–215.
- Westerbeek, H., Koolen, R., & Maes, A. (2015). Stored object knowledge and the production of referring expressions: The case of color typicality. *Frontiers in Psychology*, 6, 935.
- Wong, A. D. (2005). The reappropriation of tongzhi. *Language in Society*, 34(5), 763–793.
- Zhang, Z., Bergen, L., Paunov, A., Ryskin, R., & Gibson, E. (2023). Scalar implicature is sensitive to contextual alternatives. *Cognitive Science*, 47(2), e13238.