

This article is part of a running Theme Series entitled ‘Interaction: Talk and Beyond’, which explores non-verbal dimensions of talk in interaction. A full Series introduction can be found in issue 20/3 (2016).

Sociolinguistic interaction and identity construction: The view from game-theoretic pragmatics¹

Heather Burnett

CNRS-Université de Paris Diderot, France

Understanding the dynamics that characterize interaction between conversational participants is a fundamental goal of most theories of socially conditioned language use and identity construction through language. In this paper, I outline a class of formal tools that, I suggest, can be helpful in making progress towards this goal. More precisely, this paper explores how Bayesian signalling game models can be used to formalize key aspects of current sociolinguistic theories, and, in doing so, contribute to our knowledge of how speakers use their linguistic resources to communicate information and carve out their place in the social world. The Bayesian framework has become increasingly popular for the analysis of pragmatic phenomena of many different types, and, more generally, these models have become a dominant paradigm for the explanation of non-linguistic cognitive processes. As such, I argue that this approach has the potential to yield a formalized theory of personal and social identity construction and to situate the study of sociolinguistic interaction within a broader theory of rationalistic cognition.

Bien comprendre la dynamique qui caractérise l'interaction entre interlocuteurs dans une conversation est un des objectifs principaux de la plupart des théories traitant du comportement sociolinguistique et de la construction de l'identité par le langage. Dans cet article, je développe une classe d'outils formels qui pourrait permettre d'atteindre cet objectif. Plus précisément, cet article explore comment les jeux de signalement Bayésiens peuvent être utilisés pour formaliser les aspects centraux des théories sociolinguistiques actuelles et, ce faisant, contribuer à faire connaître comment les locuteurs utilisent leurs ressources linguistiques pour transmettre de l'information et pour creuser leur place dans la société. L'utilisation du cadre Bayésien devient de plus en plus courante pour

analyser des phénomènes pragmatiques de toutes sortes, et ces modèles constituent plus généralement un paradigme dominant pour l'explication des processus cognitifs non-linguistiques. Étant donné cela, j'argumente en faveur de l'idée que cette approche est capable de déboucher sur une théorie formalisée de la construction de l'identité personnelle et sociale, et de situer l'étude de l'interaction sociolinguistique dans une théorie plus générale de la cognition rationnelle. [French]

KEYWORDS: Pragmatics, game theory, identity construction, interaction, indexicality, Bayesian cognitive science

1. INTRODUCTION

Understanding the dynamics that characterize interaction between conversational participants is a fundamental goal of most theories of socially conditioned language use and/or identity construction through language (see Goffman 1961, 1967; Gumperz 1982a, 1982b; Bell 1984, 1997; Giles, Coupland and Coupland 1991; Ochs 1993; Eckert and McConnell-Ginet 1995; Bucholtz and Hall 2005, 2008; among many others). In this paper, I outline a class of formal tools that, I suggest, can be helpful in making progress towards this goal. More precisely, this paper explores how epistemic game theory and, more specifically, Bayesian signalling game models can be used to formalize key aspects of current sociolinguistic frameworks, and in doing so, contribute to our knowledge of how speakers use their linguistic resources to communicate information and carve out their place in the social world.

Linguistic communication and identity construction through language are extremely complex cognitive and social phenomena, and a lot of open issues in the study of language, variation and identity are very subtle. Formalization can be a very powerful tool for helping us carefully distinguish between different aspects of theoretical proposals and for precisely identifying empirical predictions made by competing analyses. This being said, in order for a mathematical approach to sociolinguistic interaction and identity construction to be helpful, we need to use a formalism that is appropriate for the type of data that we want to model, and it turns out that this is not a trivial matter. In fact, many mathematical approaches to the study of meaning allow contextual factors to play only a restricted role (see the discussions in Récanati 2004, 2010) and tend to study only the behaviour of the listener, not the speaker. As such, many formal frameworks are ill-equipped to capture the context-dependent interplay between conversational participants that lies at the heart of studies of interaction in sociolinguistics and linguistic anthropology.

Developing appropriate, mathematically precise frameworks for capturing the relation between language, meaning and use is a longstanding problem in linguistics. Already in her 1985[2011] paper, 'Feminism in linguistics', Sally

McConnell-Ginet reflects on the supposed ‘trade-off’ between formal rigour and interactivity as follows (McConnell-Ginet 2011: 64):

Many critics would say that rigor in linguistics has been achieved at the price of rigor mortis. The radical operation required to ‘isolate’ the language system has killed it: formal rules and representations provide no insight into language as a human activity. The defense against this malpractice charge, of course, is to develop an account of the relation between abstract linguistic systems and the mental states and processes, social actions and cultural values, that infuse them with life.

In this paper, I propose that game theory gives us a way to answer McConnell-Ginet’s challenge of providing ‘vibrant’ formal theories of linguistic communication. Indeed, the idea that language can be conceptualized as a game dates back at least to Wittgenstein (1953), and the proposal that concepts and mechanisms from game theory could be useful for analyzing language use has previously been explored by Goffman (1970), Bourdieu (1977), Myers-Scotton and Bolonyai (2001), Dror, Granot and Yaeger-Dror (2013, 2014), and Clark (2014), among others. In their seminal article on the potential of game theory to illuminate questions in variationist sociolinguistics Dror, Granot and Yaeger-Dror (2013: 562) say,

There is a precedent for considering a link between linguistic practices and ‘game theory’ (Benz et al., 2005). A fair amount of work has been published over recent years considering linguistic diffusion from an evolutionary game theory perspective, including work on game theory and typology (Jäger, 2007), on the semantics of numbers (Jäger, 2012) and the pragmatics of ‘epistemically lifted game’ phrasing or number choice (Franke, 2009) [...], and on Gricean theories of ‘pragmatics’ (Jäger, 2008b, 2012).

This paper picks up where Dror, Granot and Yaeger-Dror left off, detailing how (what they call) *epistemically lifted game theory* has been used in the field of formal pragmatics to formalize Gricean theories of linguistic interaction, and explores how these recent developments in pragmatics might be fruitfully applied to the modelling of sociolinguistic interaction. More technically, I focus on **signalling game architecture** (Lewis 1969) paired with a **probabilistic/Bayesian approach** to speaker/listener reasoning (see Oaksford and Chater 2007, for an overview). This framework has become increasingly popular for the analysis of pragmatic phenomena of many different types, and Bayesian game-theoretic models (more generally) have become a dominant paradigm for the explanation of non-linguistic cognitive processes (to be discussed below). As such, I argue that such models have the potential to yield a formalized theory of identity construction through language and to situate the study of sociolinguistic interaction within a broader theory of rationalistic cognition.²

In this article, I will mostly concentrate on building a model that can relate insights from formal semantics/pragmatics (and Bayesian cognitive science

more generally) to the quantitative patterns of linguistic variation that are the main object of study of quantitative/variationist sociolinguistics (Labov 1966, *et seq.*). This particular focus is motivated both by my own personal research interests and by the desire to help bridge the substantial theoretical and empirical gap that has historically existed between these areas.³ This being said, it is my hope that sociolinguists working in more qualitative approaches might also find the topics discussed here to be useful, possibly as a way to provide a precise language into which to state theories concerning the relationship between language and the interactive process. Indeed, there has been at least a small amount of interest in formal modelling in the study of interaction in the traditions of Goffman and Gumperz: in his 1970 work, *Strategic Interaction*, Goffman even develops his own class of game-theoretic models,⁴ called *expression games* (see Goffman 1970: 11–46, for definitions), for capturing linguistic (and other) behaviour associated with deception and deception detection. Gumperz (1982a) also acknowledges the potential usefulness of (appropriate) formal modelling in developing more goal-oriented interactive theories of linguistic variation. Possibly evoking evolutionary game theory (Maynard Smith and Price 1973), he says:

There is a need for a sociolinguistic theory which accounts for the communicative functions of linguistic variability and for its relation to speakers' goals without reference to untestable functionalist assumptions about conformity or nonconformance to closed systems of norms. Since speaking is interacting, such a theory must ultimately draw its basic postulates from what we know about interaction. [...] Empirical methods must be found to determine the extent to which underlying knowledge is shared – perhaps through models of social aggregates patterned on modern theories of ecosystems, which specify constraints on interpretation and behavior but do not seek to predict what is actually used and how it is evaluated. (Gumperz 1982a: 29)

Thus, I suggest that Bayesian game-theoretic modelling is one of the kinds of empirical methods that can help us discover the constraints on interpretation and behaviour that characterize situations of communicative interaction, and therefore can be of interest to researchers working in a wide range of sociolinguistic traditions.

The paper is laid out as follows: in section 2, I give a general overview of game-theoretic models and a more detailed overview of the particular class of models that are commonly employed in formal pragmatics – **signalling games**. Then I outline some recent work in game-theoretic pragmatics that formalizes Gricean reasoning, focusing on *Iterated Best Response* (Franke 2009) and *Rational Speech Act* (Frank and Goodman 2012; Goodman and Stuhlmüller 2013) models. I show how this framework has been applied to modelling the calculation of **scalar implicatures** and argue that, in these approaches, scalar reasoning shares many of the properties that have been proposed to characterize identity construction through language. These similarities, I

propose, motivate their application to sociolinguistic phenomena. In section 3, I explore such an application through giving a formalization of the *Third Wave* approach to the meaning of variation (Eckert 2000, 2008, 2012) in terms of Bayesian Signalling Game models, and I illustrate how this framework works through modelling two empirical studies: Gratton's (2016) production study of non-binary people's use of (ING) (i.e. *working* vs *workin'*); and Levon's (2014) perception study of the relationship between listener gender stereotypes and the interpretation of high/low pitch in the speech of British men. Section 4 concludes.

2. GAME THEORY AND BAYESIAN REASONING

Probably the most compelling argument in favour of using game theory to analyze the interactive and strategic aspects of sociolinguistic variation is simply that game theory is, in its essence, a mathematical formalism for describing situations of strategic interaction. In a nutshell, a **game** (as game-theorists define it) is composed of two basic parts.⁵ The first component is the **architecture** of the situation of interaction itself: for a situation to be a game (in this technical sense), there must be at least two players. The players must interact and this interaction must result in a particular outcome which must depend on the choice of strategy of each player; that is, players' actions have to play some role in determining what happens. Finally, each player must have a preferential ordering over outcomes: they must prefer some things to happen more than other things. The value that players assign to particular outcomes and actions is called their **utility**. The second crucial component to a game is the **solution concept**: a rule or algorithm that determines how the game is played, i.e. what actions the particular players take.

As stated, this definition of a game is very abstract and broad. This reflects the fact that game-theoretic models have been used to analyze very many instances of interaction (both human and non-human) across the economic, biological, social and cognitive sciences.⁶ However, when it comes to dealing with linguistic communication it makes sense to start looking at a much more narrow class of games: **signalling games** (Lewis 1969). Informally speaking,⁷ in a signalling game there are two players: the **speaker** (S) and the **listener** (L). S knows a piece of (truth conditional) information that they want to communicate to L. L wants to learn the information that S is trying to communicate to them, and in order to help them to transmit their piece of information, S has a set of messages that they can choose to send to L. In formal pragmatics, we usually assume that messages are particular linguistic forms paired with semantic meanings. S's action is to pick a message to send to L, i.e. to say something. Then L's action is to assign an interpretation to the message, i.e. to understand it in some way. The game thus has two outcomes:

1. **L interprets the message in the way that S intended**, and so they learn the information that S wanted to tell them. This outcome is good for S because they managed to communicate their information, and it is good for L because they were able to learn the information that they were looking for.
2. **L doesn't interpret the message in the way that S intended**, and so they do not learn the information. This outcome is bad for S since they did not communicate their information, and furthermore it is bad for L since they did not learn the fact that S was trying to tell them.

Since outcome 1 is preferred by both S and L (i.e. they both 'win') and outcome 2 is dispreferred by both S and L (i.e. they both 'lose'), the signalling game is a game of **cooperation** (Schelling 1960). With this architecture in mind, we now turn to the solution concept: what determines which message the speaker will pick to try to communicate their desired piece of information and which meaning the listener will assign to S's message?

When it comes to linguistic communication, be it propositional communication or identity construction, a natural idea is that both the speaker and listener's actions will be largely determined by properties of their **beliefs** about their conversational partner and their **reasoning** about how their partner will act (see the discussion in Franke 2009). A very influential recent idea in formal pragmatics is that the approach that we adopt to analyze human reasoning is the one that is found in the **Bayesian/probabilistic** approach to cognitive science (see Tenenbaum et al. 2011; Zeevat and Schmitz 2015; Franke and Jäger 2016, for recent overviews of Bayesian pragmatics). More specifically, as discussed in Tenenbaum et al. (2011: 1279), the Bayesian approach to cognitive science can be summarized as a set of answers to the following questions concerning the nature of knowledge and cognition:

- (1) a. How does abstract knowledge guide learning and inference from sparse data?
- b. What forms does abstract knowledge take, across different domains and tasks?
- c. How is abstract knowledge itself acquired?

The Bayesian answer to questions (1a) and (1c) is that learning and acquisition are products of **statistical inference**. More specifically, Bayesians propose that the fundamental rule of human reasoning is **Bayesian Inference**:⁸ humans draw a conclusion B after having observed event A (we write this as $P(B|A)$, read as *the probability of B given A*) through combining two things:

1. How likely they think A is to indicate B (written $P(A|B)$, read '*the likelihood of A given B*').
2. How likely they thought B was to begin with (written $Pr(B)$, read '*their prior belief that B is the case*').

The Bayesian answer to (1b) is that knowledge takes the form of rich, structured representations, such as phonological or syntactic tree structures (Chomsky 1957; Chomsky and Halle 1968; among many others), structured semantic representations (Link 1983; Bach 1986; among others), or even, as I will propose below, indexical fields (Eckert 2008). This approach is therefore innovative because it allows for a **synthesis** of symbolic approaches to language (common in formal linguistics) with statistical, frequency-based approaches (common in more functionally oriented linguistics). As Tenenbaum et al. (2011: 1279) say,

Until recently, cognitive modelers were forced to choose between two alternatives (Pinker, 1997): powerful statistical learning operating over the simplest, unstructured forms of knowledge [...] or richly structured symbolic knowledge equipped with only the simplest, non-statistical forms of learning, checks for logical inconsistency between hypotheses and observed data, as in nativist accounts of language acquisition (Niyogi, 2006). It appeared necessary to accept either that people's abstract knowledge is not learned or induced in a nontrivial sense from experience (hence essentially innate) or that human knowledge is not nearly as abstract or structured (as 'knowledge-like') as it seems (hence simply associations).

By virtue of their generality, Bayesian models have found wide applications across the cognitive sciences, being used to model phenomena related to vision (Kersten and Yuille 2003; Yuille and Kersten 2006; among many others), memory (Shiffrin and Steyvers 1997; Steyvers, Griffiths and Dennis 2006), sensorimotor systems (Kording and Wolpert 2006), and, of course, language (see Chater and Manning 2006; Lassiter and Goodman 2013; Zeevat and Schmitz 2015; Franke and Jäger 2016, for overviews). As such, when we propose to use Bayesian game-theoretic models to analyze identity construction, we are making a particular proposal concerning the formalization of sociolinguistic theories and, at the same time, integrating the study of the identity construction process into the broader field of Bayesian cognitive science.

In order to understand how these models work, in the next section, we will see how they can be used to formalize a particular theory of linguistic interaction, **Gricean pragmatics** (Grice 1975), and how they can be used to model a context-sensitive interactive phenomenon: *scalar implicature calculation*.

2.1 Gricean reasoning in Bayesian game-theoretic pragmatics

One of the principal pragmatic phenomena that has been treated in Bayesian game-theoretic pragmatics is **implicature calculation**: 'extra' inferences drawn by the listener that are triggered by the speaker's use of one linguistic form over another. For example, in many situations, if we hear a speaker say an

utterance with *some*, such as (2a), we will conclude the negation of the corresponding utterance with *all* (2b).

- (2) a. Mary ate **some** of the cookies.
 b. → Mary did not eat **all** of the cookies.

Although the inference (2b) may seem automatic, there are reasons to think that this **implicature** is not directly encoded into the meaning of *some*. For instance, implicature calculation is restricted to certain linguistic environments, something that would be unexpected if the *but not all* inference was entailed by the literal meaning of *some*. In particular, if *some* is embedded within the antecedent of a conditional (3a) or within a question (3b), the *but not all* implicature is not drawn, i.e. it is not possible to answer 'yes' to (3b) if you have eaten all the cookies.

- (3) a. If you eat **some** of the cookies, I'll be angry.
 → If you eat some but not all of the cookies, I'll be angry.
 b. Did you eat **some** of the cookies?
 → Did you eat some but not all of the cookies?

Instead of being an aspect of the literal meaning of *some*, meanings such as (2b) are commonly proposed to arise through the combination of the speaker's action (choosing to use *some*, rather than *all*) and the listener's particular interpretation of that choice in the discourse context (Strawson 1950; Grice 1975; Levinson 1983; Horn 1989; among very many others). Thus, at a basic level, scalar implicature have been analyzed in pragmatics along the same lines as many researchers in linguistics, anthropology and philosophy have analyzed personal and social identity: not as 'a stable structure located primarily in the individual psyche or in fixed social categories', but as 'a relational and socio-cultural phenomenon that emerges and circulates in local discourse contexts of interaction' (Bucholtz and Hall 2005: 586).

The clearest way to see how this framework works is through an example. Suppose we have two agents: the speaker (S) and the listener (L). S and L baked three cookies, and then, while L was out, Mary stopped by and possibly ate some of them. Suppose that L calls the house and wants to know how many of the cookies Mary ate. What should S say and how should L understand what S says to them?

In this example, there are four possibilities (shown in Table 1): the situation (or *world*, in formal semantics terminology) in which Mary didn't eat any of the cookies (call this w_0); the situation in which she ate one cookie (w_1); the situation in which she ate two cookies (w_2); and the situation in which she ate all three of the cookies (w_3).

Suppose that S sees that Mary actually ate two of the cookies; therefore, S wants to communicate that we are in w_2 in this example. In order to try to

Table 1: Possible worlds in the cookie example

World	Description
w_0	Mary ate 0 cookies
w_1	Mary ate 1 cookie
w_2	Mary ate 2 cookies
w_3	Mary ate 3 cookies

Table 2: Messages in the cookie example

Short name	Message	[[message]]
NONE	Mary ate none of the cookies	$\{w_0\}$
SOME	Mary ate some of the cookies	$\{w_1, w_2, w_3\}$
ALL	Mary ate all of the cookies	$\{w_3\}$

communicate this fact to L, S needs to pick a message. For simplicity, we will assume that S can choose from the three messages shown in Table 2.

For illustration, we will limit the messages in the model to those three. Of course, we could have also included the longer *Mary ate some but not all of the cookies* and the non-partitive *Mary ate some cookies* into the message set of the game. Indeed, we would include them in a comprehensive analysis of the use and interpretation of scalar quantifiers in English, but we will keep things simple here.

As shown in Table 2, messages are associated with semantic meanings, which we will take to be the sets of possible worlds/situations in which they are true. For example, *Mary ate none of the cookies* (NONE) is true only in one world: the one in which she eats zero cookies. On the other hand, *Mary ate some of the cookies* (SOME) is true in three worlds: every one except w_0 .

The speaker's first step in choosing what to say is to make a hypothesis about their interlocutor's belief state concerning which cookies may (or may not) have been eaten. Suppose that S thinks that L has no prior expectations about how many cookies Mary ate. We can represent the listener's uncertain belief state through having their prior beliefs, Pr , be **uniform** over the set of possible worlds, as shown in Table 3.

With L's prior beliefs in mind, S picks a message to say. Following Franke (2009) and Frank and Goodman (2012), when the listener hears a message m ,

Table 3: Listener (L) has uniform prior beliefs in the four possible worlds ($Pr(w)$).

Possible world	w_0	w_1	w_2	w_3
Probability	0.25	0.25	0.25	0.25

Table 4: Listener's (L) beliefs in the four possible worlds (w) immediately after hearing the message (m) ($Pr(w|m)$)*

Message	w ₀	w ₁	w ₂	w ₃
NONE	1	0	0	0
ALL	0	0	0	1
SOME	0	0.333	0.333	0.333

*Shading on the table highlights high values

the first thing that they do is restrict their attention to the worlds in which m is true. More technically, L conditions their beliefs on the meaning of the message, which is equivalent to intersection followed by renormalization of the measure. In other words, after hearing a message m , the listener 'zooms in' on the worlds in which the m is true, discards the ones in which m is false as impossible, and then re-adjusts their beliefs. Assuming L's prior beliefs are uniform (Table 3), L's beliefs immediately after hearing a message are shown in Table 4: after hearing NONE, L is certain that Mary ate zero cookies; after hearing ALL, L is certain that she ate the three cookies; and after hearing SOME, L is certain that Mary did not eat zero cookies, but is equally uncertain about how many she ate.

Following Grice, we assume that speakers aim to make the most informative statement possible, and informativity supplies a point for S and L to coordinate on (Lewis 1969).⁹ These models formalize Grice's maxims of quantity¹⁰ by building informativity into the speaker's *utility function* (U_S). S's utility function is a measure of how useful a message would be for S to communicate their desired piece of information to L. Following Frank and Goodman (2012), who follow Shannon (1948), the informativity of a message is measured as the natural log of the listener's beliefs conditioned on the meaning of the message, as shown in (4).

(4) Utility of a message m to communicate w ($U_S(m, w)$):

$$U_S(m, w) = \ln(Pr(w|m))$$

In signalling games, speaker utility functions often also encode information associated with **costs** for different messages. For example, if we were also considering messages like *Mary ate some but not all the cookies*, we might want to assign it a penalty to reflect the fact that it is much longer than *Mary ate some of the cookies*. Likewise, if we are comparing *Mary ate some of the cookies* with *Mary ate some cookies*, we might want to penalize the non-partitive sentence, since listeners have been shown to be less likely to draw implicatures in the partitive construction than in the non-partitive (Degen 2015). Since, in our small example, there is no major length or other grammatical difference between the messages under consideration, we will not incorporate message

Table 5: Speaker's utility for message (m) for communicating world (w) ($U_S(m, w)$)*

Message	w_0	w_1	w_2	w_3
NONE	0	$-\infty$	$-\infty$	$-\infty$
ALL	$-\infty$	$-\infty$	$-\infty$	0
SOME	$-\infty$	-0.108	-0.108	-0.108

*values have been rounded to three decimal places, where appropriate; shading highlights high values

costs into the speaker utility function in this paper; however, message costs are an important way in which linguistic conditioning factors can be captured in game-theoretic pragmatics.

Thus, to generate measures of utility for each message, we plug the numbers in Table 4 into the equation in (4), which generates Table 5. Messages whose denotation does not contain a world are assigned the utility $-\infty$ for communicating that world because $\ln(0) = -\infty$. So, by virtue of the fact that the conditionalized probability of interpreting w_0 after immediately hearing *SOME* is zero (see Table 4), the speaker's utility of using *SOME* to communicate w_0 is as low as it possibly can be: $-\infty$. As shown in Table 5, the only useful message for communicating w_0 is *NONE* (since $\ln(1) = 0$); the only useful message for communicating w_1 and w_2 is *SOME* (since $\ln(0.33\bar{3}) \approx -0.108$); and the most useful message for communicating that Mary ate the three cookies is *ALL* (because $0 > -0.108$).

One of the great benefits of Bayesian game-theoretic models is that they can be used to make gradient quantitative predictions concerning linguistic production and interpretation. This arises under the hypothesis that speakers are **approximately rational**: they are trying to make the choice that will have the best chance of accomplishing their goals (whatever they may be) (see Anderson 1991; among others), but they may not always pick the optimal action. That is to say, if human speakers were fully rational, we would always expect them to perform the action that has the highest utility. So, if *S* wishes to communicate that we are in w_1 or w_2 , we would expect them to pick *SOME* 100 percent of the time; whereas, if they wish to communicate w_3 , then we expect them to pick *ALL* 100 percent of the time (never *SOME*). However, we know that human mental computation can be impeded by a variety of time/resource constraints (fatigue, working memory, etc.), and, as discussed above, scalar implicature generated and calculated has been observed to be variable. So to account for variability in action selection, we assume that the speaker chooses which message to say non-deterministically: using the *Soft-Max* choice rule (Luce 1959; Sutton and Barto 1998).¹¹ With this rule, the probability of making a choice increases as the utility of the choice increases, and the amount of variability is governed by the values of a parameter α , called the *temperature*. A parameter in a model is a particular number that is chosen for a

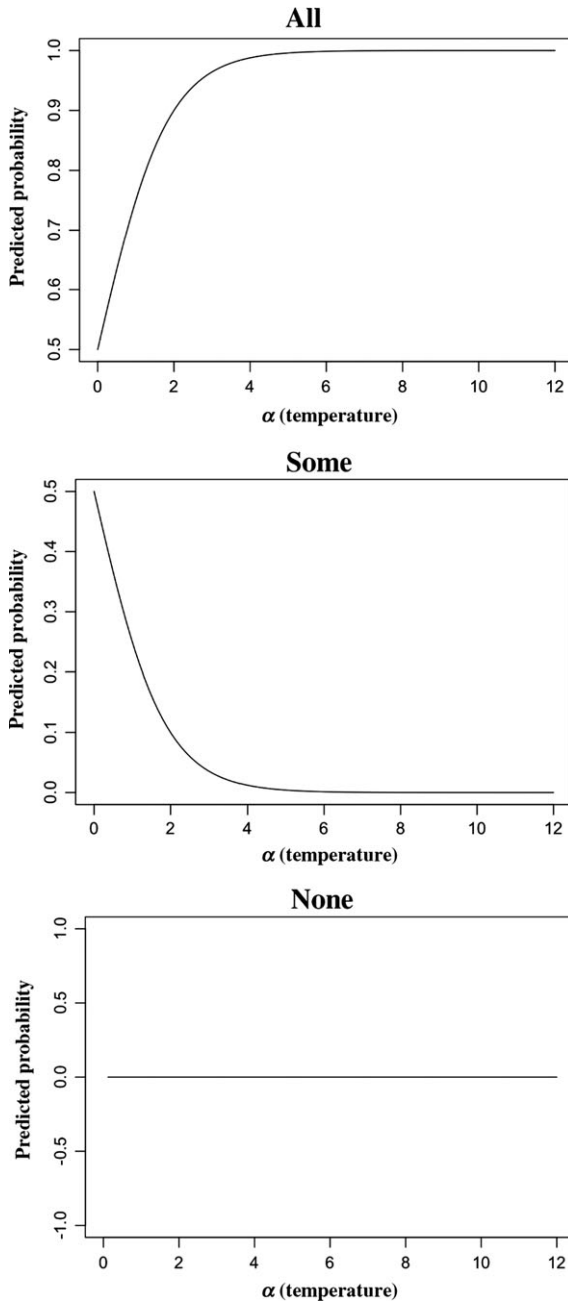


Figure 1: Predicted probabilities of using ALL VS SOME VS NONE to communicate w_3 , varying α

particular dataset, which contributes to determining the model's predictions for that dataset.

In this context, we use α to encode how much inherent variability we think there is in the system that we are studying. When α is set to ∞ (infinity), the speaker picks the message with the highest utility 100 percent of the time (so no inherent variability); whereas, anything lower than ∞ will predict some variability, even if it is just a tiny amount. The lower the value of α is set, the more variable the speaker's choice will be. The basic structure of the model therefore predicts a **range** of probabilities for the use of a variant; whereas, exact probabilities will depend on which number for α is selected. For example, the predicted probabilities of using NONE, SOME VS ALL to communicate that we are in w_3 for values of α less than 12 are shown in Figure 1. Note that once α gets higher than around six, choice becomes almost deterministic in this model. In particular, as α gets larger than zero, the probability of using ALL sharply rises and the probability of using SOME sharply decreases. NONE is never predicted to be used.

When we are modelling actual quantitative studies, the value for α that best fits the observed data can be estimated (as in Goodman and Stuhlmüller 2013; Franke and Degen 2016); indeed, in the next section, we will choose values for α that best fit the quantitative sociolinguistic variation data that we will try to capture. However, in order to exemplify how the scalar implicature model works, I will pick a value for α that will allow for some variation: $\alpha = 10$. Plugging the utility values in Table 5 and $\alpha = 10$ into the Soft-Max choice rule generates the probability distributions over messages shown in Table 6. If S wants to communicate w_0 , then the model predicts that they will categorically say *Mary ate none of the cookies*. If S wants to communicate w_1 or w_2 , then the model predicts that they will categorically say *Mary ate some of the cookies*. However, if the speaker wants to communicate w_3 , the model predicts variable language use: S will say *Mary ate all of the cookies* 99 percent of the time, but *Mary ate some of the cookies* a highly disfavoured one percent of the time.

Finally, listeners interpret messages using their hypotheses that speakers are (approximately) rational and motivated by informativity, combined with their prior beliefs. In other words, as discussed in the previous section, we treat linguistic interpretation as Bayesian inference, and, from the values in Table 6

Table 6: Speaker's predicted use of m , given w with $\alpha = 10$ ($P_S(m|w)$)*

Message	w_0	w_1	w_2	w_3
NONE	1	0	0	0
ALL	0	0	0	0.99
SOME	0	1	1	0.01
Prediction:	Cat. NONE	Cat. SOME	Cat. SOME	Favoured ALL

*cat.=categorically; shading highlights high values

(the probability that S would use m given w) and the values in Table 3 (L's prior beliefs concerning w), we derive a probability distribution over interpretations of messages, as shown in Table 7.¹² Crucially, Table 7 should be read from the listener's perspective, inverted from Table 6 which is from the speaker's perspective. For example, Table 6 indicates that the probability of S using ALL to communicate w_3 is 0.99; whereas, the same cell in Table 7 indicates that L has a probability of 1 (100%) of interpreting w_3 if they hear ALL.

In particular, the model predicts that if L hears NONE, then they will categorically understand that Mary ate zero cookies. However, if L hears SOME, then we predict that they will largely hesitate between w_1 and w_2 , with a very tiny probability assigned to w_3 . In other words, the model predicts that if the speaker says *Mary ate some of the cookies*, the listener will almost certainly understand that Mary ate some but not all of the cookies. Hence, the model predicts the (variable) *but not all* implicature.

Note that the exact probability distributions shown above are dependent on parameters of the model. In particular, for exposition, we made the assumption that L started off with no prior expectation concerning how many cookies Mary would eat. However, suppose that L knows that Mary usually likes to have two cookies for her dessert. So, before hearing what S has to say, they are expecting Mary to have eaten two cookies. This belief can be represented through changing L's prior probability distribution to one that is heavily weighted on w_2 , as shown in Table 8.

In this case, L's interpretation probabilities change, and L's probability of interpreting w_2 after hearing SOME jumps up to 0.87 from 0.498. This is because, as discussed above, Bayesian interpretation takes L's prior beliefs into account. Since L's prior probability distribution assigns w_2 a probability of 0.7 and w_1 only a probability of 0.1, L will think that it is much more likely that S is trying to tell them that we are in w_2 than in w_1 if they hear SOME. In other words, in these models, aspects of listener's prior beliefs and expectations can

Table 7: Listener's predicted interpretation of w , given m ($P_L(w|m)$)*

Message	w_0	w_1	w_2	w_3	Prediction
NONE	1	0	0	0	Categorical w_0
ALL	0	0	0	1	Categorical w_3
SOME	0	0.498	0.498	0.005	Favoured w_1, w_2

*values have been rounded to three decimal places; shading highlights high values

Table 8: Listener's prior probability distribution heavily weighted on w_2

Possible world	w_0	w_1	w_2	w_3
Probability	0.1	0.1	0.7	0.1

dramatically influence how they interpret the linguistic expressions that S offers to them.

So far in this paper, we have seen that Bayesian game-theoretic models provide a framework for capturing instances of interactive co-construction of propositional meaning, as exemplified by scalar implicature calculation. Observe that, in this framework, the *not all* inference is not part of the semantic meaning of *SOME*. Rather, this implicature arises as a product of coordination between the speaker and listener based on reasoning about a set of **scalar alternatives** (*NONE VS SOME VS ALL*). In the rest of the article, we will explore how Bayesian game-theoretic models can be applied to capturing the linguistic co-construction of personal and social identity as a product of speaker-listener coordination based on reasoning about other kinds of alternatives: **sociolinguistic variants**.

3. IDENTITY CONSTRUCTION IN GAME-THEORETIC PRAGMATICS

This section presents a new formal model of identity co-construction based on the Bayesian signalling game models commonly used in game-theoretic pragmatics, defined above. I define the models and then show how they can be applied to the analysis of two empirical studies of sociolinguistic variation and interpretation. But first, I identify a number of empirical generalizations taken from the literature on sociolinguistic perception/interpretation that, I argue, a formal model of socially conditioned variation and identity construction through language should capture.

3.1 *Social meaning, variation and identity construction*

One of the most common ways that social meaning and sociolinguistic interpretation has been investigated is through the use of a particular experimental paradigm called the *Matched Guise Technique* (MGT; Lambert et al. 1960). This is an experimental method widely used in social psychology and more recently in variationist sociolinguistics to assess listeners' implicit attitudes towards speakers of different linguistic varieties. In this paradigm, participants listen to samples of recorded speech that have been designed to differ in specific and controlled ways. Each participant hears one of two recordings of the same speaker (called *guises*) which differ only in the alternation studied. In this way, we know that any significant differences between guises that we may observe are due to the linguistic forms under study, and not to something else (content, other aspects of voices, etc.). After hearing a recording, participants' attitudes towards the recorded speaker are assessed, usually via focus groups and/or questionnaires.

In a series of papers on the social interpretation of variable (ING) (5), Campbell-Kibler (2006, 2007, 2008) reports on a MGT study that she

performed with 124 participants using stimuli formed from the speech of eight different speakers.

- (5) a. I'm **working** on my paper. -*ing*
 b. I'm **workin'** on my paper. -*in'*

Given its size, this study yielded a variety of complex patterns. Among them was that all speakers were rated as significantly more *educated* and more *articulate* in their *-ing* guises than in their *-in'* guises. Likewise, speakers were significantly more likely to be described as a *redneck* in their *-in'* guises than in their *-ing* guises. Finally, one male speaker, who she calls *Jason*, is significantly more likely to be described as *gay* when he says *-ing* than when he says *-in'*.

Similarly, Podesva et al. (2015) performed a MGT study with 70 participants investigating the interpretation of /t/ release (6) using stimuli formed from political speeches of six American politicians (Barak Obama, John Edwards, Nancy Pelosi, George W. Bush, Hillary Clinton, and Condoleezza Rice).

- (6) /t/ release
 a. wa[t^h]er released /t/
 b. wa[r]er flapped /t/

As in Campbell-Kibler's study, the /t/ release study yielded a number of results concerning associations with released vs flapped/unreleased /t/: for example, John Edwards and Condoleezza Rice were rated as significantly more *articulate* in their released /t/ guises than in their flapped guise. On the other hand, Nancy Pelosi was rated as significantly less *friendly* and less *sincere* when she used released /t/. Thus, the speaker's choice of linguistic form can often make a large difference in how they are perceived by the listener.

This being said, which exact property attributions a particular variant will trigger often depends on which other properties the listener believes are held by the speaker. This can already be seen in Campbell-Kibler and Podesva et al.'s MGT studies: Podesva et al. found significant relationships between articulateness and released /t/ with Edwards and Rice, but only with these two speakers. Likewise, flapping/not releasing made only Nancy Pelosi sound more friendly and sincere. Finally, in Campbell-Kibler (2007), only Jason was heard as gay when he said *-ing*; (ING) made no differences to interpretations of sexual orientation with the other speakers. Thus, listener prior beliefs concerning individual speakers appear to constrain sociolinguistic interpretation.

Furthermore, we have evidence that listener prior beliefs concerning groups of speakers, i.e. *stereotypes*, also influence interpretation. This can be clearly seen in Levon (2014) which reports on a MGT study of the interpretation of styles involving higher ([+raised]) vs lower ([-raised]) pitch in the speech of British men, among other variables. In addition to the standard MGT response questionnaire, Levon had participants fill out the *Male Role Attitudes Survey* (MRAS; Pleck, Sonenstein and Ku 1994). This instrument measures agreement with a series of statements corresponding to 'traditional' male

gender norms. In other studies, higher scores on the MRAS have been shown to correlate with higher rates of homophobia, church attendance, promiscuity and other behaviours. Again, Levon's (2014) study generated a lot of different empirical results; however, the one which is pertinent for this paper is that participants separated into two groups (Burnett and Levon 2016): listeners scoring high on the MRAS attributed more **incompetent** and less **masculine** personae to speakers using [+raised] pitch than [-raised] pitch, and listeners scoring low on the MRAS attributed more incompetent personae to speakers using [+raised] pitch, but there was no difference in masculinity. Thus, listeners' prior ideological beliefs concerning the relationship between masculinity and (in)competence play a role in determining what kinds of personae/identities they attribute to speakers based on their linguistic performances.

In summary, we see from the literature on sociolinguistic perception that hearers make judgments about the properties that characterize speakers based on the linguistic forms that they use; however, social interpretation is crucially constrained by listener prior beliefs. This being said, interpretation is only one side of the coin when it comes to sociolinguistic interaction, and we have reason to believe that speakers strategically exploit listeners' interpretation processes to communicate properties about themselves and, in doing so, construct their identities.

A clear illustration of this phenomenon comes from Kiesling's (1998) study of the use of (ING) by nine college fraternity members. Among other contexts, Kiesling made recordings of the boys socializing and participating in an organizational meeting. He found that many of the fraternity members displayed a higher rate of *-in'* in the informal socializing context (75% *-in'* on average) than in the more formal meeting context (47% *-in'*; Kiesling 1998: 76). For instance, one of the boys, Mack, uses the *-in'* form 73 percent of the time in socialization; however, in the meeting his rate of *-in'* drops to 13 percent. This being said, some fraternity members do not decrease their use of *-in'* in the meeting context like Mack does: another boy, Speed, moves only from 95 percent *-in'* to 82 percent, a non-significant difference. According to Kiesling, both Mack's change across contexts and Mack and Speed's divergences are due to the fact that the different fraternity boys are aiming to construct different identities in different contexts. While both boys adopt a casual persona in socialization, Speed and Mack are constructing different kinds of powerful masculine personae at the meeting. Following Ochs (1992), among others, Kiesling proposes that the *-in'* form indexes (i.e. is related to) properties such as *casualness*, *physical masculinity* (see Connell 2005) and an *anti-establishment* stance; whereas, *-ing* indexes *formality*, *non-physical masculinity/non-masculinity* and a *pro-establishment* stance. Based on analysis of Speed's speech, Kiesling proposes that Speed's identity at the meeting is 'a rebel creating a powerful stance vis-a-vis the ability power hierarchy in the fraternity that rewards hard work, rather than structural power for its own

sake' (Kiesling 1998: 86), which explains why Speed favours the *-in'* form in this context. Mack, on the other hand, aims to construct the persona of the 'leader who knows what is good for his flock' (Kiesling 1998: 92). Kiesling argues that there is 'strong evidence that Mack, who used less [*-in'*] in the meeting than while socializing, is indeed displaying a meeting identity by indexing structurally powerful alignment roles, roles that are also associated with the use of the [*-ing*] variant' (Kiesling 1998: 92).

Finally, we also have reasons to believe that speakers are sensitive to what their interlocutors think about them, and that these expectations can influence which linguistic forms they use. This can be seen clearly in Gratton's (2016) recent study of the link between (ING) and gender presentation across contexts. Gratton conducted group interviews with non-binary individuals; that is, individuals whose gender identity does not respect the male/female binary. She focused on two people: Flynn, who was assigned female at birth, and Casey, who was assigned male at birth. Gratton conducted two sets of interviews with these consultants: the first one was in a queer-friendly environment (their home and a queer coffee shop, respectively) and the second one took place in a public coffee shop. Gratton found that neither Flynn nor Casey show a significant difference in their use of *-in'* vs *-ing* in the queer-friendly environments: 44 and 58 percent *-ing* respectively. However, in the public coffee shop, Flynn, who was assigned female at birth, uses significantly more *-in'* (80%), while Casey, who was assigned male at birth, uses significantly more *-ing* (89%). Based on ethnographic analysis of her interviews, Gratton argues that these patterns are created by variation in how her speakers evaluate their listeners' beliefs. She says:

The individuals in this community believe that in queer environments, they can be read as non-normative quite easily, which means, according to them, that they do not need to consciously worry about their gender presentation. However, the same cannot be said for non-queer public contexts. They believe that individuals whom they encounter in non-queer public spaces will pre-suppose a binary gender based mainly on their physiological characteristics. In order to present a non-binary or non-normative gender identity, they must distance themselves from the gender which is presumed – their gender-assigned-at-birth – by utilizing resources which resist the gender norms associated with their respective gender-assigned-at-birth. (Gratton 2016: 6)

In summary, in this section, I outlined a number of properties that, based on research in sociolinguistics, a formal model of identity construction should capture. I argued that we want a framework that can predict variable, quantitative patterns of variation and interpretation, and that can model interaction between conversational participants such that the speaker (tries to) choose the variant that has the best chance to construct their desired persona. Furthermore, we want a model in which the listener's prior beliefs and ideologies constrain interpretation. I highlight that these are many of the same

properties that we saw characterize scalar implicature calculation. Therefore, I propose that it is reasonable to extend the game-theoretic models outlined in the previous section to model the strategic aspect of sociolinguistic variation.

3.2 Formalization of the Third Wave

In the construction of our formal model, we will build on the large amount of previous work theorizing about identity construction through language in sociolinguistics, linguistic anthropology and sociocultural linguistics more generally (see Bucholtz and Hall 2005, 2008). In particular, we will adopt the *indexicality* theory of social meaning in which abstract mental representations mediate the relationship between language and identity categories (Silverstein 1976, 1979, 2003; Ochs 1992, 1993; Eckert 2008; among others), and, more specifically, how indexicality is developed and related to patterns of language use within the *Third Wave* approach to variation (see Eckert 2012, for an overview).

In the model, as in a classic signalling game, there are two players: the speaker (S) and the listener (L). Third Wave variation theory focuses on how variants combine together to form styles, which construct particular identities or personae (see Podesva 2004; Eckert 2008; Zhang 2008; among many others). In this paper, we will take personae to be particular collections of properties that 'go together'. For our empirical illustration, we will construct personae from the set of properties in (7), where we understand *masculine/feminine* to broadly regroup various kinds of masculinities/non-femininities and non-masculinities/femininities respectively (see Cameron and Kulick 2003; Eckert and McConnell-Ginet 2013; among others).

(7) {competent, incompetent, casual, delicate, masculine, feminine}

We also suppose that there are some ideological relations (for more formal definitions, see Burnett 2017) between these properties – for example, suppose that one cannot be both *competent* and *incompetent* at the same time; nor can one be *casual* and *delicate* at the same time, or *masculine* and *feminine* at the same time. So, given this setup, the set of personae generated from (7) is shown in (8).

(8) Set of personae

1. {competent, casual, masculine}
2. {competent, casual, feminine}
3. {competent, delicate, masculine}
4. {competent, delicate, feminine}
5. {incompetent, casual, masculine}
6. {incompetent, casual, feminine}
7. {incompetent, delicate, masculine}
8. {incompetent, delicate, feminine}

(ING) and gender identity construction. As in signalling games, we will have a set of messages that the speaker can pick from to try to construct their desired personae. As an illustration, we will first consider Gratton's (2016) study of (ING) and gender identity construction across contexts. The messages for this game are shown in (9).

(9) Messages = {-ing, -in'}

In Third Wave variation theory, individual variants have meaning that goes beyond their truth conditional meaning. In particular, variants index sets of properties/stances, called their *indexical field* (Eckert 2008). Following (simplified) Eckert (2008), we propose the indexical fields for (ING) shown in Table 9.

The representations in Table 9 show the most standard approach to the representation of indexical fields. This being said, the 'Eckert' fields shown will not be exactly the right kinds of objects to be integrated into the larger model. Instead, in the spirit of Montague (1973), we will adopt an **equivalent** characterization indexical fields as the set of personae that they have the potential to construct. The relation between more traditional Eckert fields and equivalent Eckert-Montague fields is shown in Table 10. Observe that, although both variants have the potential to construct competent and casual personae (for example), only *-ing* can construct competent and delicate personae, while only *-in'* can construct incompetent and casual personae.

As in the scalar implicature example, the speaker makes a hypothesis about the listener's prior beliefs concerning which persona(e) they instantiate.¹³

Table 9: Eckert indexical fields for (ING)

Variant	Eckert field
<i>-ing</i>	{competent, delicate}
<i>-in'</i>	{incompetent, casual}

Table 10: Messages and indexical fields

Variant	Eckert field	Eckert-Montague field
<i>-ing</i>	{competent, delicate}	{comp., delicate, masc.}, {comp., delicate, fem.}, {comp., casual, masc.}, {comp., casual, fem.}, {incomp., delicate, masc.}, {incomp., delicate, fem.}
<i>-in'</i>	{incompetent, casual}	{incomp., casual, masc.}, {incomp., casual, fem.}, {comp., casual, masc.}, {comp., casual, fem.}, {incomp., delicate, masc.}, {incomp. delicate, fem.}

*comp.=competent; incomp.=incompetent; fem.=feminine; masc.=masculine

Thus, we represent the listener's prior beliefs as a probability distribution over personae (*Pr*). *Pr* can encode specific beliefs about the particular speaker (*Mary has X properties*) or more general stereotypes (*Canadian women have X properties*). For example, suppose listeners think that *delicateness* and *femininity* are ideologically linked (Ochs 1992), as are *casualness* and *masculinity*. This can be represented in the listener's prior beliefs as there being much higher probability mass on personae that have the combination {casual, masculine} and {delicate, feminine} than on {casual, feminine} and {delicate, masculine}. A prior distribution that encodes these ideological relations is shown in Table 11: the persona {competent, casual, masculine} has a higher weight than the persona {competent, delicate, masculine} (0.1 vs 0.05), and the same holds for {incompetent, casual, masculine} compared to {incompetent, delicate, masculine}.

In her study of Flynn (the non-binary individual who was assigned female at birth), Gratton (2016) argues that, in the public coffee shop, they are worried that their interlocutor will incorrectly attribute them a feminine persona. We also represent this specific belief in the model through putting more probability mass on feminine personae than on masculine personae in the listener's prior beliefs, as shown in Table 11. In other words, if we add up the probabilities on feminine personae in Table 11, we see that Flynn thinks that their coffee shop interlocutor assigns them 0.7 probability of being some feminine persona.

As in the models for implicature calculation presented in section 2, we assume that the first thing that the listener does when they hear a variant is to restrict their attention to the personae that appear in the variants' Eckert-Montague field, and adjust their prior beliefs accordingly. For example, if they hear the variant *-ing*, they discard the possibility that the speaker is incompetent and casual (i.e. they assign 0 probability to {incompetent, casual, masculine} and {incompetent, casual, feminine} personae). Likewise, if they hear *-in'*, they assign 0 probability to {competent, delicate, feminine} and {competent, delicate, masculine} personae.

Table 11: Listener's prior beliefs (Flynn/public coffee shop)*

Persona	<i>Pr</i> (persona)
{competent, casual, masculine}	0.1
{competent, casual, feminine}	0.1
{competent, delicate, masculine}	0.05
{competent, delicate, feminine}	0.25
{incompetent, casual, masculine}	0.1
{incompetent, casual, feminine}	0.1
{incompetent, delicate, masculine}	0.05
{incompetent, delicate, feminine}	0.25

*Shading highlights high values

In parallel with the scalar implicature model, we assume that speaker utility is guided by informativity, and, for simplicity, we will assume that there are no cost differences between variants.¹⁴ So, transposed into the identity construction model, speaker utility (U_S) for a variant m to construct a persona P is given by (10): the utility for the speaker to use a variant m , given that they wish to construct persona P , is the natural log of the probability of P conditioned on (i.e. taking into account) the indexical fields of m .

$$(10) U_S(P, m) = \ln(P|m)$$

Following Gratton, we assume that Flynn wants to construct a non-feminine persona (for example, a {competent, casual, masculine} persona). So in exactly the same way that we calculated utilities for *SOME*, *NONE* and *ALL* in the previous section, we plug the values in Table 11 into the speaker utility equation in (10), which assigns a utility of -1.946 for Flynn to use *-in'* and a utility of -2.079 to use *-ing*. We can then take these utilities and plug them into the Soft-Max choice rule described in the previous section. If we do so, and we set the temperature $\alpha = 10$ for this dataset, we predict that they will use *-in'* 79 percent of the time, which is similar to what Gratton found.¹⁵

Thus, given a particular proposal concerning a speaker's analysis of their communication situation (who they want to be and who others think they are), we predict which variant the speaker will be most likely to use. Note, crucially, that the framework does not assume that all or even most aspects of message/interpretation selection or utility calculation are conscious or intentional. Indeed, as discussed in section 2, one of the core proposals of this paper is that identity construction arises from the same principles and mechanisms as other kinds of cognitive activities. As mentioned above, Bayesian models have been shown to be useful in the analysis of cognitive processes like vision, perception, memory and motion planning, which are all activities that feature very little conscious awareness. So it by no means follows that because it is appropriate to model aspects of cognition as instances of rationalistic decision making, all such decision making must be operating above the level of consciousness (see Dennett [1993]; Graziano [2013] for general discussion and Burnett [2017] for more discussion of this point in the context of sociolinguistic variation and identity construction).

For Casey, who was assigned male at birth, the listener prior beliefs look different: Casey is worried about being attributed a masculine persona in the coffee shop. So, although the general ideological relationships between masculinity, casualness, femininity and delicateness stay the same in Table 12, Casey's representation of their listener's prior beliefs about them is highly weighted on masculine personae.

Based on these priors, if we suppose that Casey wishes to construct a feminine persona (for example {competent, delicate, feminine}), the model predicts that they will categorically use *-ing*.¹⁶ So given different listener prior

Table 12: Listener's prior beliefs (Casey/public coffee shop)*

Persona	Pr(persona)
{competent, casual, masculine}	0.25
{competent, casual, feminine}	0.05
{competent, delicate, masculine}	0.1
{competent, delicate, feminine}	0.1
{incompetent, casual, masculine}	0.25
{incompetent, casual, feminine}	0.05
{incompetent, delicate, masculine}	0.1
{incompetent, delicate, feminine}	0.1

*Shading highlights high values

Table 13: Listener's prior beliefs (Flynn/queer friendly environment)*

Persona	Pr(persona)
{competent, casual, masculine}	0.79
{competent, casual, feminine}	0.01
{competent, delicate, masculine}	0.01
{competent, delicate, feminine}	0.05
{incompetent, casual, masculine}	0.05
{incompetent, casual, feminine}	0.01
{incompetent, delicate, masculine}	0.01
{incompetent, delicate, feminine}	0.05

*Shading highlights high values

beliefs and different desired identities, we predict drastically different patterns of variation.

Of course, speakers' assessments of how they are perceived can change in different situations. Suppose, as described by Gratton, they are fairly certain that they will be attributed their desired persona ({competent, casual, masculine} for Flynn) when they are in a queer friendly environment (Table 13). In this case, particular choices of variants will have less of an effect, and the model predicts that Flynn will use *-in'* 50 percent of the time (rather than 79% in the public setting), with $\alpha = 10$.

More generally, the predicted probabilities for Flynn's language use in the public vs queer-friendly contexts are depicted in Figure 2. This figure shows the model's predicted probabilities with all values of the parameter α less than 50. Recall that lower values for α encode more inherent variability in the system, while higher values make the speaker's choice more deterministic. Thus, the models capture the observation that speakers' different

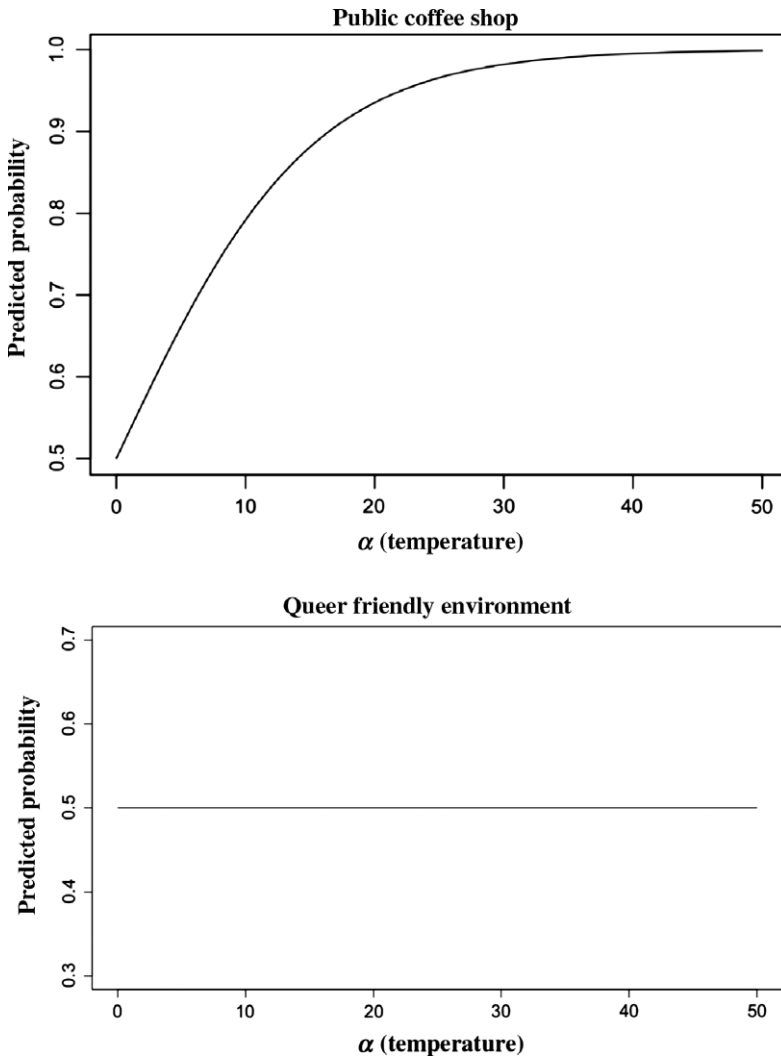


Figure 2: Predicted probabilities for Flynn's use of *-in'* across contexts for $\alpha \leq 50$

representations of listeners' beliefs can have a large effect on their language use.

Up till now, I have been arguing that Bayesian game-theoretic models of the type we have been discussing can formally express the link between observed patterns of language use and certain proposals within sociolinguistics and linguistic anthropology (such as those of Ochs, Eckert, Campbell-Kibler and Gratton), which previously had only been stated at an intuitive level. However, in addition to helping us 'pin down' some of the insights of existing proposals, I

Table 14: Eckert indexical fields for (ING) (competing analysis)

Variant	Eckert field
<i>-ing</i>	{competent}
<i>-in'</i>	{incompetent}

argue that game-theoretic modelling can also be a tool for **testing** hypotheses concerning difficult theoretical questions like: (a) which indexical fields to assign to which variants; and (b) what speaker ideologies look like. For example, we might wonder whether there are other analyses of the indexical fields of (ING) besides the one proposed in Table 9 that are consistent with Gratton's results. For instance, perhaps the proposal in Table 9 is unnecessarily complex and we could do just as well with the simpler analysis in Table 14 where the more 'standard' form *-ing* indexes *competence* and the 'vernacular' form *-in'* indexes *incompetence*.

In this case, if we keep the same set up as above (i.e. we suppose that Flynn's desired persona is {competent, casual, masculine} and that their assumed listener priors are as in Table 11), we make the prediction that Flynn will use *-ing* **100 percent** of the time under all values of α , which is clearly very different from what Gratton found. Thus, Bayesian game-theoretic models provide a way of constructing direct **empirical** arguments in favour of one analysis of the indexical field over another.

Pitch and gender identity construction. Different listener prior beliefs can also produce different sociolinguistic interpretations in the model. To see this, we can consider the Levon (2014) example discussed in section 3.1. Recall that Levon found a distinction between two groups of listeners: *progressives*, who had no particular (diagnosable) stereotypes concerning male behavior; and *conservatives*, whose beliefs conformed to particular 'traditional' male gender norms. Furthermore, Levon found that differences in male stereotypes translated into differences in interpretation of pitch: conservatives assigned incompetent and non-masculine personae to speakers using raised pitch; whereas, progressives assigned incompetent and both masculine and non-masculine personae to raised pitch (Burnett and Levon 2016).

This pattern is predicted by the model: in this game, we have two messages (11), and I propose that [+pitch raised] indexes incompetence, while [-pitch raised] indexes competence (Table 15).

(11) Messages = {[+pitch raised], [-pitch raised]}

Given Levon's result associated with the MRAS, it is reasonable to suppose that a main difference between progressives and conservatives lies in which kinds of personae they expect to encounter in the world: progressives in

Table 15: Messages and indexical fields (Levon 2014)*

Variant	Eckert field	Eckert-Montague field
[+pitch raised]	{incompetent}	{incomp., delicate, masc.}, {incomp., delicate, fem.}, {incomp., casual, masc.}, {incomp., casual, fem.}
[-pitch raised]	{competent}	{comp., casual, masc.}, {comp., casual, fem.}, {comp., delicate, masc.}, {comp., delicate, fem.}

*comp.=competent; incomp.=incompetent; fem.=feminine; masc.=masculine

Table 16: Conservatives' prior beliefs ($Pr(\text{persona})$) in Levon (2014)*

Persona	$Pr(\text{persona})$
{competent, casual, masculine}	0.15
{competent, casual, feminine}	0.15
{competent, delicate, masculine}	0.15
{competent, delicate, feminine}	0.15
{incompetent, casual, masculine}	0.05
{incompetent, casual, feminine}	0.15
{incompetent, delicate, masculine}	0.05
{incompetent, delicate, feminine}	0.15

*Shading highlights high values

Levon's study have no beliefs about how incompetence and non-masculinity/femininity cluster together; therefore, we represent their beliefs as a uniform probability distribution over personae. On the other hand, conservatives in Levon's study think that incompetence is emasculating. We represent this belief in the model as one that it is highly unlikely that they will encounter an incompetent, masculine person (see Table 16).

Based on these prior beliefs (with α again set at 10), the model correctly predicts that speakers using [+pitch raised] will be more likely to be attributed non-masculine personae by conservative listeners; whereas, it is predicted that there will be no such gender-based asymmetry for progressives, as shown in Table 17.

Thus, these formal models can also be useful for capturing the relationship between ideological structure and sociolinguistic interpretation.

4. CONCLUSION

In this paper, I gave an exploration of how formal tools commonly used in the growing field of game-theoretic pragmatics could be applied to the analysis of

Table 17: Probability distribution over interpretations [+ pitch raised]*

Persona	Progressives	Conservatives
{competent, casual, masculine}	0	0
{competent, casual, feminine}	0	0
{competent, delicate, masculine}	0	0
{competent, delicate, feminine}	0	0
{incompetent, casual, masculine}	0.25	0.13
{incompetent, casual, feminine}	0.25	0.37
{incompetent, delicate, masculine}	0.25	0.13
{incompetent, delicate, feminine}	0.25	0.37

*Shading highlights high values

socially conditioned variable linguistic phenomena and identity construction through language. Using a Bayesian game-theoretic framework inspired by the *Iterated Best Response/Rational Speech Act* models, I showed how we can formalize sociolinguistic theories, in this case, *Third Wave variation* theory, in order to make both qualitative and quantitative predictions about variable language use and interpretation across different contexts and different kinds of speakers. I argued that, unlike many existing mathematical semantic/pragmatic frameworks, these models can capture the interactive co-construction of meaning that forms the basis of how we establish our place in the social world. In particular, inferences associated with properties of the speaker, such as (12b), are not proposed to be directly encoded into the message, but rather arise as the product of (un)conscious coordination between the speaker and the listener.

- (12) a. I have been workin' on my paper.
 b. → The speaker is casual.

Furthermore, I argued that the proposed models can capture the contribution that speaker/listener specific prior beliefs and ideologies make to social interpretation and patterns of context-sensitive linguistic variation. I therefore propose that these models constitute yet another item in the sociolinguist's 'toolkit' for analyzing patterns of variation and interaction.

Importantly, I stress that game-theoretic modelling **complements** existing methodologies in sociolinguistics rather than replacing any aspect of current practice: logistic regression/Goldvarb analyses are crucial for telling us which linguistic and social factors are operative in the dataset that we are studying. In other words, variationist analyses help us identify what the empirical generalizations that our analyses should capture are; whereas, game-theoretic analyses allow us to state different (possibly competing) analyses in such a way that allows us to automatically determine whether or not they do in fact

capture these generalizations. More specifically, the framework that I have developed provides a way to relate analyses consisting of indexical fields and ideological structure to probability distributions over variants and interpretations. However, this architecture remains an empty shell in the absence of concrete detailed proposals concerning which properties are in the indexical fields of which variants, and which ideological relations exist between properties. Qualitative ethnographic analysis allows us to identify (or at least to make reasonable hypotheses concerning) the ideological structure assumed by the speakers in the communities under study and to formulate analyses which can then be evaluated using the models. In this way, I suggest that game-theoretic modelling can serve as a **link** between proposals concerning social meaning and ideological structure, which have been greatly studied by researchers in interactional sociolinguistics and conversational analysis, and the fine-grained patterns of grammatical variation studied in more quantitative approaches.

Furthermore, the shape of this link has its source in the principles underlying Bayesian cognitive science, which have independently been argued to characterize numerous aspects of human cognitive activities. Consequently, we might also expect sociolinguistic studies to have an important role to play in the future development of this exciting new field. I therefore conclude that, working in tandem with detailed quantitative and ethnographic studies of linguistic interaction, Bayesian game-theoretic models have useful applications to treat patterns of sociolinguistic variation/interpretation and the potential to situate the study of sociolinguistic interaction 'beyond talk', i.e. within a broader theory of human behaviour and cognition.

NOTES

1. This research has been partially supported by the program 'Investissements d'Avenir' overseen by the French National Research Agency, ANR-10-LABX-0083 (Labex EFL), and a fellowship from the Center for the Study of Language and Information at Stanford University. I thank Leon Bergen, Adrian Brasoveanu, Judith Degen, Fabio del Prete, Bernard Fradin, Chantal Gratton, Erez Levon, Eric McCreedy, Jessica Rett, Devyani Sharma, Elizabeth Smith, Sali Tagliamonte, Meredith Tamminga, audiences at UCL, Institut Jean Nicod, LLF Paris-Diderot, Stanford, UCLA, UCSC and NWAV45, and especially Penny Eckert and Dan Lassiter for very helpful comments and discussions. All errors are my own.
2. The focus of this work will be the construction of a mathematical model of identity construction and its place within the cognitive sciences. Readers interested in the relationship between social meaning and other kinds of pragmatic meaning within the game-theoretic perspective are referred to Burnett (2017).

3. Fortunately, this gap is starting to close with recent works such as Acton and Potts (2014), Acton (2014), Beltrama (2016), Burnett (2017), and this work aims to contribute to this research programme in 'socio-semantics'.
4. *Strategic Interaction* was written after Goffman spent a sabbatical year at Harvard working with the Nobel-prize-winning game theorist Thomas Schelling (Manning 1992). Thus, in addition to a contribution to the study of the dynamics of deception, Goffman's expression games also make a contribution to the field of game theory proper.
5. For a more technical introduction, see Osborne and Rubinstein (1994). For linguistically oriented introductions, see Benz, Jäger and van Rooij (2005), Jäger (2011) (game theory+pragmatics), or Dror, Granot and Yaeger-Dror (2013, 2014) (game theory+sociolinguistics).
6. In addition to the cognitive processes discussed in this paper, game-theoretic analyses have been applied to an enormous range of topics from human and animal population dynamics and reproduction, to voting strategy, auctions and pricing, and economic and political bargaining.
7. Readers who are more formally inclined may enjoy the linguistic and philosophical introduction to signalling games in Franke (2009).
8. This inference is laid out in its gory detail here: $(B_i \setminus A) = \frac{\Pr(B_i) \times P(A|B_i)}{\sum_{j=1}^{|B|} \Pr(B_j) \times P(A|B_j)}$
9. Note that, under this view, technically speaking, the listener does not have to be positively disposed towards the speaker or actively wish to engage in some meaning making process with them; they just try to extract the most information from S's utterance as possible.
10. (Grice 1975: 45): (1) make your contribution as informative as is required (for the current purposes of the exchange); (2) do not make your contribution more informative than is required.
11. The Soft-Max choice rule is written out formally as follows: for a world w , a message m and a real number α , called the *temperature*,

$$P_s(m|w) = \frac{\exp(\alpha \times U_s(w,m))}{\sum_{m' \in M} \exp(\alpha \times U_s(w,m'))}$$
12. Calculating these probability distributions by hand can be tedious. Fortunately, to facilitate calculations and prediction testing, a number of computational implementations of the Rational Speech Act model have been developed that the interested reader might find helpful:
 - a. Potts' implementation in python: <https://github.com/cgpotts/pypragmods>
 - b. Goodman and Tenenbaum's implementation in Church: <https://probmods.org/>
 - c. Goodman and Stuhlmüller's implementation in WebPPL: <http://dippl.org/examples/pragmatics.html>
13. This can be seen as a way of encoding of *audience design* (Bell 1984) into the model.
14. Again we could also incorporate message costs into the model as a way of capturing both social and linguistic conditioning factors. For example, we know that (ING) is conditioned by grammatical category and other abstract properties of morphological structure (Labov 1966; Houston 1985; Tamminga 2014). However, in order to do this properly, we would need more complicated

message representations. So the unification of social and linguistic factors within Bayesian game-theoretic pragmatics is left to future research.

15. These calculations can, again, be tedious to do by hand. Therefore, a computational implementation in python of the model used in this paper (done in collaboration with D. Lassiter) is available for interested readers at <https://github.com/hsburnett/smg>
16. In order to capture variation in Casey's use of (ING) in the coffee shop, we could attribute Casey's *-in'* to linguistic/grammatical conditioning and/or adopt a more complicated model which allows some variation in persona selection. See Burnett (2017) for such a model.

REFERENCES

- Acton, Eric K. 2014. Pragmatics and the social meaning of determiners. Unpublished PhD dissertation. Stanford, California: Stanford University.
- Acton, Eric K. and Christopher Potts. 2014. That straight talk: Sarah Palin and the sociolinguistics of demonstratives. *Journal of Sociolinguistics* 18: 3–31.
- Anderson, John R. 1991. Is human cognition adaptive? *Behavioral and Brain Sciences* 14: 471–485.
- Bach, Emmon. 1986. The algebra of events. *Linguistics and Philosophy* 9: 5–16.
- Bell, Allan. 1984. Language style as audience design. *Language in Society* 13: 145–204.
- Bell, Allan. 1997. Language style as audience design. In Nikolas Coupland and Adam Jaworski (eds.) *Sociolinguistics: A Reader and Coursebook*. New York: St Martin's Press. 240–250.
- Beltrama, Andrea. 2016. Bridging the gap. Intensifiers between semantic and social meaning. Unpublished PhD thesis. Chicago, Illinois: University of Chicago.
- Benz, Anton, Gerhard Jäger and van Rooij Robert (eds.). 2005. *Game Theory and Pragmatics*. London: Palgrave Macmillan UK.
- Bourdieu, Pierre. 1977. The economics of linguistic exchanges. *Social Science Information* 16: 645–668.
- Bucholtz, Mary and Kira Hall. 2005. Identity and interaction: A sociocultural linguistic approach. *Discourse Studies* 7: 585–614.
- Bucholtz, Mary and Kira Hall. 2008. All of the above: New coalitions in sociocultural linguistics1. *Journal of Sociolinguistics* 12: 401–431.
- Burnett, Heather. 2017. Signalling games, sociolinguistic variation and the construction of style. Accepted in *Linguistics & Philosophy*.
- Burnett, Heather and Erez Levon. 2016. A (more or less) compositional semantics of style. Paper presented at Meaning, Optimization and Interaction (MOI) workshop, 9 September, Université Paris-Diderot, France.
- Cameron, Deborah and Don Kulick. 2003. *Language and Sexuality*. Cambridge, U.K.: Cambridge University Press.
- Campbell-Kibler, Kathryn. 2006. Listener perceptions of sociolinguistic variables: The case of (ING). Unpublished PhD thesis. Stanford, California: Stanford University.
- Campbell-Kibler, Kathryn. 2007. Accent, (ING), and the social logic of listener perceptions. *American Speech* 82: 32–64.
- Campbell-Kibler, Kathryn. 2008. I'll be the judge of that: Diversity in social perceptions of (ING). *Language in Society* 37: 637–659.

- Chater, Nick and Christopher D. Manning. 2006. Probabilistic models of language processing and acquisition. *Trends in Cognitive Science* 10: 335–344.
- Chomsky, Noam. 1957. *Syntactic Structures*. The Hague, The Netherlands: Mouton.
- Chomsky, Noam and Morris Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.
- Clark, Robin L. 2014. *Meaningful Games: Exploring Language with Game Theory*. London: MIT Press.
- Connell, Raewyn W. 2005. *Masculinities*. Berkeley, California: University of California Press.
- Degen, Judith. 2015. Investigating the distribution of some (but not all) implicatures using corpora and web-based methods. *Semantics and Pragmatics* 8(11): 1–55.
- Dennett, Daniel C. 1993. *Consciousness Explained*. London: Penguin U.K.
- Dror, Moshe, Daniel Granot and Malcah Yaeger-Dror. 2013. Speech variation, utility, and game theory. *Language and Linguistics Compass* 7: 561–579.
- Dror, Moshe, Daniel Granot and Malcah Yaeger-Dror. 2014. Teaching and learning guide for speech variation, utility, and game theory. *Language and Linguistics Compass* 8: 230–242.
- Eckert, Penelope. 2000. *Language Variation as Social Practice: The Linguistic Construction of Identity in Belten High*. Oxford, U.K.: Wiley-Blackwell.
- Eckert, Penelope. 2008. Variation and the indexical field. *Journal of Sociolinguistics* 12: 453–476.
- Eckert, Penelope. 2012. Three waves of variation study: The emergence of meaning in the study of sociolinguistic variation. *Annual Review of Anthropology* 41: 87–100.
- Eckert, Penelope and Sally McConnell-Ginet. 1995. Constructing meaning, constructing selves. In Kira Hall and Mary Bucholtz (eds.) *Gender Articulated: Language and the Socially Constructed Self*. New York: Routledge. 469–507.
- Eckert, Penelope and Sally McConnell-Ginet. 2013. *Language and Gender*. Cambridge, U.K.: Cambridge University Press.
- Frank, Michael C. and Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336: 998–998.
- Franke, Michael. 2009. Signal to act: Game theory in pragmatics. Unpublished PhD thesis. Tübingen, Germany: University of Tübingen.
- Franke, Michael and Judith Degen. 2016. Reasoning in reference games: Individual- vs. population-level probabilistic modeling. *PloS One* 11: e0154854.
- Franke, Michael and Gerhard Jäger. 2016. Probabilistic pragmatics, or why Bayes' rule is probably important for pragmatics. *Zeitschrift für Sprachwissenschaft* 35: 3–44.
- Giles, Howard, Nikolas Coupland and Justine Coupland. 1991. Accommodation theory: Communication, context, and consequence. In Howard Giles, Nikolas Coupland and Justine Coupland (eds.) *Contexts of Accommodation* (Developments in Applied Sociolinguistics, 1). Cambridge, U.K.: Cambridge University Press. 1–68.
- Goffman, Erving. 1961. *Encounters: Two Studies in the Sociology of Interaction*. Indianapolis, Indiana: Bobbs-Merrill.
- Goffman, Erving. 1967. *Interaction Ritual: Essays on Face-to-Face Interaction*. Chicago, Illinois: Aldine.
- Goffman, Erving. 1970. *Strategic Interaction* (Volume 1). Philadelphia, Pennsylvania: University of Pennsylvania Press.
- Goodman, Noah D. and Andreas Stuhlmüller. 2013. Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science* 5: 173–184.
- Gratton, Chantal. 2016. Resisting the gender binary: The use of (ING) in the construction of non-binary transgender identities. *Penn Working Papers in Linguistics* 22(2): article 7.

- Graziano, Michael S. 2013. *Consciousness and the Social Brain*. Oxford, U.K.: Oxford University Press.
- Grice, H. Paul. 1975. Logic and conversation. *Syntax and Semantics* 3: 41–58.
- Gumperz, John J. 1982a. *Discourse Strategies* (Volume 1). Cambridge, U.K.: Cambridge University Press.
- Gumperz, John J. 1982b. *Language and Social Identity* (Volume 2). Cambridge, U.K.: Cambridge University Press.
- Horn, Laurence R. 1989. *A Natural History of Negation*. Chicago, Illinois: University of Chicago Press.
- Houston, Ann Celeste. 1985. Continuity and change in English morphology: The variable (ING). Unpublished PhD thesis. Philadelphia, Pennsylvania: University of Pennsylvania.
- Jäger, Gerhard. 2011. Game-theoretical pragmatics. In Johan van Benthem and Alice ter Meulen (eds.) *Handbook of Logic and Language*. Amsterdam, The Netherlands: Elsevier. 467–491.
- Kersten, Daniel and Alan Yuille. 2003. Bayesian models of object perception. *Current Opinion in Neurobiology* 13: 150–158.
- Kiesling, Scott Fabius. 1998. Men's identities and sociolinguistic variation: The case of fraternity men. *Journal of Sociolinguistics* 2: 69–99.
- Kording, Konrad P. and Daniel M. Wolpert. 2006. Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences* 10: 319–326.
- Labov, William. 1966. *The Social Stratification of English in New York City*. Washington, D.C.: Center for Applied Linguistics.
- Lambert, Wallace E., Richard C. Hodgson, Robert C. Gardner and Samuel Fillenbaum. 1960. Evaluational reactions to spoken languages. *The Journal of Abnormal and Social Psychology* 60: 44–51.
- Lassiter, Daniel and Noah D. Goodman. 2013. Context, scale structure, and statistics in the interpretation of positive-form adjectives. *Semantics and Linguistic Theory* 23: 587–610.
- Levinson, Stephen C. 1983. *Pragmatics*. Cambridge, U.K.: Cambridge University Press.
- Levon, Erez. 2014. Categories, stereotypes, and the linguistic perception of sexuality. *Language in Society* 43: 539–566.
- Lewis, David. 1969. *Convention*. Cambridge, Massachusetts: Harvard University Press.
- Link, Godehard. 1983. The logical analysis of plurals and mass terms. In Rainer Bäuerle, Christoph Schwarze and Arnim von Stechow (eds.) *Meaning, Use, and Interpretation of Language*. Berlin, Germany/New York: DeGruyter. 302–323.
- Luce, R. Duncan. 1959. On the possible psychophysical laws. *Psychological Review* 66: 81–95.
- Manning, Philip. 1992. *Erving Goffman and Modern Sociology*. Cambridge, U.K.: Polity.
- Maynard Smith, J. and G. Price. 1973. The logic of animal conflict. *Nature* 246: 15–18.
- McConnell-Ginet, Sally. 2011. *Gender, Sexuality and Meaning: Linguistic Practice and Politics*. Oxford, U.K.: Oxford University Press.
- Montague, Richard. 1973. The proper treatment of quantification in ordinary English. In K. J. Hintikka, J. M. E. Moravcsik and P. Suppes (eds.) *Approaches to Natural Language* (Proceedings of the 1970 Stanford Workshop on Grammar and Semantics). Dordrecht, Holland: D. Reidel Publishing Company. 221–242.
- Myers-Scotton, Carol and Agnes Bolonyai. 2001. Calculating speakers: Codeswitching in a rational choice model. *Language in Society* 30: 1–28.
- Niyogi, Partha. 2006. *The Computational Nature of Language Learning and Evolution*. Cambridge, Massachusetts: MIT Press.

- Oaksford, Mike and Nick Chater. 2007. *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford, U.K.: Oxford University Press.
- Ochs, Elinor. 1992. Indexing gender. *Rethinking Context: Language as an Interactive Phenomenon* 11: 335–358.
- Ochs, Elinor. 1993. Constructing social identity: A language socialization perspective. *Research on Language and Social Interaction* 26: 287–306.
- Osborne, Martin J. and Ariel Rubinstein. 1994. *A Course in Game Theory*. Cambridge, Massachusetts: MIT Press.
- Pleck, Joseph H., Freya L. Sonenstein and Leighton C. Ku. 1994. Attitudes toward male roles among adolescent males: A discriminant validity analysis. *Sex Roles* 30: 481–501.
- Podesva, Robert. 2004. On constructing social meaning with stop release bursts. Paper presented at Sociolinguistics Symposium 15, 1–5 April, Newcastle upon Tyne, U.K.
- Podesva, Robert J., Jermy Reynolds, Patrick Callier and Jessica Baptiste. 2015. Constraints on the social meaning of released/t: A production and perception study of US politicians. *Language Variation and Change* 27: 59–87.
- Recanati, François. 2004. *Literal Meaning*. Cambridge, U.K.: Cambridge University Press.
- Recanati, François. 2010. *Truth-Conditional Pragmatics*. Oxford, U.K.: Clarendon Press.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, Massachusetts: Harvard University Press.
- Shannon, Claude. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27: 379–423.
- Shiffrin, Richard M. and Mark Steyvers. 1997. A model for recognition memory: REM – retrieving effectively from memory. *Psychonomic Bulletin & Review* 4: 145–166.
- Silverstein, Michael. 1976. Shifters, linguistic categories, and cultural description. *Meaning in Anthropology* 1: 1–55.
- Silverstein, Michael. 1979. Language structure and linguistic ideology. In Paul Clyne, William F. Hanks and Carol F. Hofbauer (eds.) *The Elements: A Parasession on Linguistic Units and Levels* (Papers from the Conference on Non-Slavic Languages of the USSR). Chicago, Illinois: Chicago Linguistic Society. 193–247.
- Silverstein, Michael. 2003. Indexical order and the dialectics of sociolinguistic life. *Language & Communication* 23: 193–229.
- Smith, E. Allyn, Kathleen Hall and Benjamin Munson. 2010. Bringing semantics to sociophonetics: Social variables and secondary entailments. *Laboratory Phonology* 1: 121–155.
- Steyvers, Mark, Thomas Griffiths and Simon Dennis. 2006. Probabilistic inference in human semantic memory. *Trends in Cognitive Science* 10: 327–334.
- Strawson, P. F. 1950. On referring. *Mind* 59: 320–344.
- Sutton, Richard S. and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. Cambridge, Massachusetts: MIT press.
- Tamminga, Meredith. 2014. Persistence in the production of linguistic variation. Unpublished PhD thesis. Philadelphia, Pennsylvania: University of Pennsylvania.
- Tenenbaum, Joshua B., Charles Kemp, Thomas L. Griffiths and Noah D. Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science* 331: 1279–1285.
- Wittgenstein, Ludwig. 1953. *Philosophical Investigations [Philosophische Untersuchungen]*. London: Macmillan.
- Yuille, Alan and Daniel Kersten. 2006. Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Science* 10: 301–308.

- Zeevat, Henk and Hans-Christian Schmitz. 2015. *Bayesian Natural Language Semantics and Pragmatics* (Volume 2). Heidelberg, Switzerland: Springer.
- Zhang, Qing. 2008. Rhotacization and the 'Beijing Smooth Operator': The social meaning of a linguistic variable. *Journal of Sociolinguistics* 12: 201–222.
-

Address correspondence to:

Heather Burnett
Laboratoire de Linguistique Formelle
UMR7110 CNRS-Université de Paris Diderot
Case Postale 7031 – 5, rue Thomas Mann
F-75205 Paris Cedex 13 France
heather.susan.burnett@gmail.com