Game Theoretic Modelling for Social Meaning and Semantic Change

Heather Burnett E. Allyn Smith

Laboratoire de Linguistique Formelle
CNRS-Paris-Diderot
CSLI, Stanford University

Département de linguistique
Université du Québec à Montréal

NWAV45 (SFU-UVic)

Goal

To introduce a new class of formal tools that can be used to give a formal treatment of social factors that condition language variation and change.

Hypothesis

Bayesian signalling game models in particular have the potential to help us study how we use our linguistic resources to communicate information and carve out our place in the social world.

What is Game Theory?

Game theory is a mathematical formalism for describing situations of strategic interaction.

- *Since at least the late 1970s it has been possible to say with confidence that game theory is the most important and useful tool in the analyst's kit whenever she confronts situations in which what counts as one agent's best action (for her) depends on expectations about what one or more other agents [might] do, and what counts as their best actions (for them) similarly depend on expectations about her.*

Ross, Don. (2016, Winter update) "Game Theory." *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.) http://plato.stanford.edu/entries/game-theory/.

In what ways could Game Theory be a useful tool for sociolinguistics?

- A lot of issues of meaning and variation are very subtle and the various options can often be difficult to distinguish based on our intuitive understanding of the theoretical proposals that we are working with.
- A formal framework forces us to specify exactly what we think down to the last gory detail.
- Our different formalizations will tell us where the important difference(s) between theoretical proposals lie.

Condition

We need models that actually make empirical predictions, and we need a formalism that is appropriate for the type of data we want to model.

Why might Game Theory be better than other formal approaches to meaning on the market?

- *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.*

Sally McConnell-Ginet. (1985). “Feminism in Linguistics." *For Alma Mater: Theory and Practice in Feminist Scholarship*. Ed. P. Treichler et al.. U of Illinois P, 159-76.

What makes Game Theory appropriate for modelling sociolinguistic data?

Restoring interactivity to formal models

- We want to model the context-dependent co-creation of meaning between conversational participants found in all kinds of meaning-making, including identity construction and diachronic change.
- Though each participant makes her own choices, the interplay between them can shift meanings for both.
- This is exactly where game theory applies: situations where the outcome (interpretation) of a participant's performance depends on other participants!

Perspective of this workshop

Given that game theoretic models already exist for some kinds of pragmatic phenomena, we ask:

What if social meaning and pragmatic meaning involve the same inference-making process?

There exist fundamental similarities between (1) and (2).

(1) I'm work[in] on my paper.
$\leadsto$ The speaker is friendly. Social implicature

(2) Mary ate some of the cookies.
$\leadsto$ Mary ate some but not all of the cookies.
Scalar implicature

Plan

Game Theory Basics

Basic components of a game (Jäger, 2011)

- There are (at least) two players.
- The players interact and the interaction results in a particular outcome.
- The outcome of the interaction depends on the choice of strategy of each player.
- Each player has a preference ordering over outcomes.
    - Preferences are usually encoded as numerical values (utilities or payoffs) that are assigned to possible outcomes.

Signalling Games (Lewis, 1969)

In a nutshell. . .

A game of coordination between two agents (S (speaker/sender), L (listener/receiver)).

- S has a piece of information that L does not have (their *type*, which for our purposes is usually a meaning that S would like to communicate).
- S's action is to choose a message m to send L.
- L's action is to assign an interpretation to m (making a guess about S's type).
- S and L's payoffs are calculated based on coordination.
  - Usually, both players win if L correctly interprets S's message.
  - Typically, they both lose if S's type and L's interpretation of her type do not converge.

Example 1: Scalar Implicatures

(3) Porkchop ate some of the chocolates.
$\rightsquigarrow$ Porkchop ate some but not all of the chocolates.

How do we model the fact that we usually interpret ‘some’ as ‘some but not all’?

- Players: Serena (S) and Lenny (L).
- Set of types that Serena has (i.e. the real-life situations she might have knowledge of):
  - P ate all of the chocolates
  - P ate none of the chocolates
  - P ate some number between none and all of the chocolates.

Example 1: Scalar Implicatures

- Set of messages that Serena can send to convey her type/meaning:
  - *Porkchop ate all of the chocolates*
  - *Porkchop ate some of the chocolates*
  - *Porkchop ate some but not all of the chocolates*
  - *Porkchop didn't eat any of the chocolates*

- The costs associated with various messages. Perhaps the third message is costly because it is longer than the second. Perhaps the fourth message is costly because it includes negation, which is harder to process.

- Interpretations of these messages (actions by the receiver that guess at the type): same set as the types (P ate all, none, or an in-between number of chocolates)

Example 1: Scalar Implicatures

Preferences and Payoffs:

- Serena's preferences: be honest, avoid costly messages, try to have her type understood by Lenny
- Lenny's preferences: assume Serena is being honest, try to retrieve the real state of affairs (Serena's type)
- Note that some of these preferences may be more important than others. Serena may care more about honesty than sending a long message, for example.
- Payoff structure: all of the possible combinations of types, messages and responses and their benefit to both participants based on their preferences.

Example 1: Scalar Implicatures

Payoff/Utility:

- If the real type (state of affairs) is that P ate between none and all of the chocolates, (and this is known by S and not L), we could have:
  - S says *Porkchop ate some of the chocolates* and L understands 'P ate between none and all choc.' : (3,3)
  - S says *Porkchop ate some but not all of the chocolates* and L understands 'P ate between none and all choc.' : (2,3)
  - S says *Porkchop ate some of the chocolates* and L understands 'P ate all of the chocolates' : (0,0)
  - S says *Porkchop ate all of the chocolates* and L understands 'P ate between none and all choc.' : (-2,1)
  - Etc.

Example 2: Social Implicatures

(4) I'm work[in] on my paper.
$\rightsquigarrow$ The speaker is friendly, incompetent.
$\rightsquigarrow$ The speaker is friendly, competent.
$\rightsquigarrow$ The speaker is aloof, incompetent.

(5) I'm work[ing] on my paper.
$\rightsquigarrow$ The speaker is aloof, competent.
$\rightsquigarrow$ The speaker is aloof, incompetent.
$\rightsquigarrow$ The speaker is friendly, competent.

Set of types for Serena:

Persona	Nickname
{competent, friendly}	'cool guy/gal'
{competent, aloof}	'serious student'
{incompetent, friendly}	'doofus'
{incomptent, aloof}	'arrogant asshole'

Example 2: Social Implicatures

- Set of messages that Serena can send to convey her type/meaning:
  - *I'm workin' on my paper.*
  - *I'm working on my paper.*
- The costs associated with various messages. At least in the case of the first two (that interest us the most), there is no clear discernible cost difference in terms of production length, etc.
- Interpretations of these messages (actions by receiver that guess at the type): same set as the types (cool gal, serious student, doofus, arrogant asshole)

Example 2: Social Implicatures

Preferences and Payoffs are related in that S and L might have different preferences in different contexts, and those will in turn affect the payoffs. For example:

Context 1:

- Serena's preferences: seem friendly and competent (cool gal)
- Lenny's preferences: assess Serena's type
- Payoff structure:
  - S says *workin'* and L thinks S is a cool gal: (2,2)
  - S says *workin'* and L thinks S is a doofus: (1,1)
  - S says *workin'* and L thinks S is an arrogant ass: (0,0)
  - S says *workin'* and L thinks S is a serious student: (1,1)
  - Etc.

Example 2: Social Implicatures

Context 2:

- Serena's preferences: seem friendly and incompetent (doofus)
- Lenny's preferences: assess Serena's type
- Payoff structure:
  - S says *workin'* and L thinks S is a doofus: (2,2)
  - S says *workin'* and L thinks S is a cool gal: (1,1)
  - S says *workin'* and L thinks S is an arrogant ass: (1,1)
  - S says *workin'* and L thinks S is a serious student: (0,0)
  - Etc.

Example 3: Grammaticalization

(6) Jane is sorting the mail. (Progressive)

(7) Jane sorts the mail. (Imperfective)

As noted by Deo (2015), inter alia, progressives often become imperfectives over time, and then new progressive markers arise. How do we model this cyclic change?

- Players: Serena (S) and Lenny (L).
- Set of types that Serena has (i.e. the real-life situations she is trying to convey):
  - Something happening right now (answering *Why is Porkchop wagging his tail?*)
  - Something structural about the world (answering *Why does a dog wag his tail?*

Example 3: Grammaticalization

- Set of messages that Serena can send to convey her type/meaning:
  - *He is lying in wait of the mailman.* (PROG)
  - *He lies in wait of the mailman.* (IMPF)
- The costs associated with various messages. Having multiple forms in a similar semantic domain is costly, so having a choice of message rather than one increases the cost in the system.
- Interpretations of these messages (actions by receiver that guess at the type): same set as the types (happening now or happens generally)

Example 3: Grammaticalization

Preferences and Payoffs:

- Serena's preferences: avoid costly messages, try to have her type understood by Lenny
- Lenny's preferences: try to coordinate with Serena
- Payoff/Utility Structure: For a game of change over time, the setup will differ depending on which stage in the cycle we are in. The number of messages will also change, which affects payoffs.
- In this case, we are at a stage where IMPF and PROG both exist but where IMPF can be used to respond to both questions, while PROG can only respond to 'now' questions. Thus, we are moving toward a system where IMPF can only be used for 'general' questions.

Example 3: Grammaticalization

Payoff/Utility example:

- If the real type (state of affairs) is conveying that something is happening right now:
  - S says *He is lying in wait of the mailman.* and L understands ‘now’ : (3,3)
  - S says *He is lying in wait of the mailman* and L understands ‘general’ : (0,-2)
  - S says *He lies in wait of the mailman.* and L understands ‘now’ : (2,3)
  - S says *He lies in wait of the mailman.* and L understands ‘general’ : (0,0)

Gamify your research

In small groups with those around you

What kinds of empirical phenomena would you like to model? What patterns in your own data would you like to represent formally to make further predictions? How would these translate into a game setup?

Consider:

- The set of types that S has knowledge of
- The set of messages that S can send
- The costs associated with various messages
- The possible interpretations of the various messages
- The players' preferences and their relative importance
- The payoff structure for various outcomes for each player

Bayesian Game-Theoretic Pragmatics

- When modelling communication, the solution concept that we use should make reference to reasoning process of the agents involved (Franke, 2009).

Hypothesis

Agents' reasoning is Bayesian (Tenenbaum et al., 2011, see)[for overview].

- Wide applications across cognitive science: perception (Yuille and Kersten, 2006), memory (Shiffrin and Steyvers, 1997), sensorimotor systems (Steyvers et al., 2006), and language (Chater and Manning, 2006).

Bayesian inference

Humans draw a conclusion B after having observed event A ($P(B|A)$) through combining:

1. How likely they think A is to indicate B (P(A|B)).
2. How likely they thought B was to begin with (Pr(B)).

Bayes rule

(8) $$P(B_i|A) = \frac{Pr(B_i) \times P(A|B_i)}{\sum_{j=1}^{|B|} Pr(B_j) \times P(A|B_j)}$$

(9) $$P(B|A) \propto Pr(B) \times P(A|B)$$

Scalar Implicatures

(10) Mary ate some of the cookies.
$\rightsquigarrow$ Mary ate some but not all of the cookies.

- We want a model that can take into account the literal meaning of a sentence.

(11)
a. Mary ate some of the cookies; in fact, she ate all of them.
b. Mary ate some of the cookies; #in fact, she ate none of them.

- Scalar enrichment is variable (Sperber and Wilson, 1986; Levinson, 2000; Degen, 2015; Degen and Tanenhaus, 2015).
  - We want a model that generates variable output.

Iterated Best Response/Rational Speech Act

- A family of similar approaches (Franke, 2009; Frank and Goodman, 2012; Goodman and Stuhlmüller, 2013; Lassiter and Goodman, 2015; Degen et al., 2015; Bergen et al., 2016; Franke and Jäger, 2016, among many others).
- A formalization of Gricean reasoning (Grice, 1975), particularly quantity and quality.

(12) Maxims of Quantity

a. Make your contribution as informative as is required (for the current purposes of the exchange).

b. Do not make your contribution more informative than is required.

(13) Maxim of Quality

Try to make your contribution one that is true.

The scenario

Suppose we baked three cookies. And L wants to know how many of them Mary ate.

Possible World (W)	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: Universe in cookie example

Suppose S sees that Mary ate two of them.

- So S wants to tell L that we are in w_2.

Messages

S can choose between three messages.

Short name	m	$[\![m]\!]$
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	$\{w_3\}$

Table: Messages in cookie example

What does the speaker do?

S makes a hypothesis about L's prior beliefs concerning which cookies may (or may not) have been eaten: *Pr*.

- Suppose S thinks L doesn't have any opinion about the cookies.

w_0	w_1	w_2	w_3
0.25	0.25	0.25	0.25

Table: L has uniform priors (Pr(w)).

Formalization of Quality Maxim

When they hear a message *m*, L restricts their attention to the worlds in which *m* is true.

- L conditions on $[\![m]\!]$: intersection followed by renormalization of the measure.

Message	w_0	w_1	w_2	w_3
NONE	1	0	0	0
ALL	0	0	0	1
SOME	0	$0.\dot{3}$	$0.\dot{3}$	$0.\dot{3}$

Table: L's beliefs immediately after hearing m (Pr(w|m)).

Formalization of Quantity Maxim

Coordination (i.e. communication) occurs because speakers try to say the most informative statement possible. And listeners know this.

- Informativity is encoded as part of S's utility function (U_S).
- Frank and Goodman (2012) following Shannon (1948): informativity is encoded as natural log of $Pr(w|m)$.

(14) $U_S(m,w) = ln(Pr(w|m)) - Cost(m)$

Costs as linguistic factors

Costs can encode grammatical/psychological constraints on utterances.

- We will ignore $C(m)$ in this quick demonstration.

Speaker Utility

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

Table: S's utility for m for communicating w (U_S(w, m)).

Predicting linguistic production

Hypothesis: Agents are approximately rational

1. Rationality: They are trying to maximize utility.
2. Approximately: They may not always pick the optimal action.
   - Computation can be impeded by time/resource constraints.

To account for variability in action selection:

Soft-Max Choice (Luce, 1959; Sutton and Barto, 1998)

For a world w, a message m and a value α (the temperature).

$$P_S(m|w) = \frac{exp(\alpha \times U_S(w, m))}{\sum_{m' \in M} exp(\alpha \times U_S(w, m'))}$$

- α introduces some non-determinancy into the model.

Quantitative Predictions for Language Use

Message	w_0	w_1	w_2	w_3
NONE	1	0	0	0
ALL	0	0	0	0.9
SOME	0	1	1	0.1
Prediction	Cat. NONE	Cat. SOME	Cat. SOME	Favored ALL

Table: S's predicted use of m, given w with $\alpha = 2$ ($P_S(\mathrm{m}|\mathrm{w})$).

Quantitative Predictions for Language Interpretation

Interpretation as Bayesian Inference

Listeners interpret messages using their hypotheses that speakers (approximately) rational and motivated by informativity, combined with their prior beliefs.

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	0.49	0.49	0.01	Favoured w_1, w_2

Table: L's predicted interpretation of w, given m (P_L(w|m)).

Heavily Weighted Priors

- Suppose L knows that Mary usually likes to have two cookies for her dessert...

w_0	w_1	w_2	w_3
0.1	0.1	0.7	0.1

Table: L's prior heavily weighted on w_2.

Prior beliefs influence interpretation

L's interpretation probabilities change.

- L's probability of interpreting w_2 after SOME is now 0.79.

Computational resources for Bayesian pragmatics

To facilitate calculations and prediction testing, a number of computational implementations have been developed:

1. Chris Potts' implementation in python:
   `https://github.com/cgpotts/pypragmods`
2. Goodman and Tenenbaum's implementation in Church:
   `https://probmods.org/`.
   - Also comes with a textbook.
3. Goodman and Stuhlmüller's implementation in WebPPL:
   `http://dippl.org/examples/pragmatics.html`
   - Also comes with a textbook for Scontras & Tessler's 2016 ESSLLI course:
     `http://gscontras.github.io/ESSLLI-2016/`

http://gscontras.github.io/ESSLLI-2016/

Modeling pragmatic inference

Day 2: Enriching the literal interpretations

Application 1: Scalar implicature

Scalar implicature stands as the poster child of pragmatic inference. Utterances are strengthened—via implicature—from a relatively weak literal interpretation to a pragmatic interpretation that goes beyond the literal semantics: "Some of the apples are red," an utterance compatible with all of the apples being red, gets strengthed to "Some but not all of the apples are red." The mechanisms underlying this process have been discussed at length. Goodman and Stuhlmüller (2013) apply an RSA treatment to the phenomenon and formally articulate the model by which scalar implicatures get calculated.

Assume a world with three apples; zero, one, two, or three of those apples may be red:

```
// possible states of the world
var statePrior = function() {
  return uniformDraw([0, 1, 2, 3])
};
statePrior() // sample a state
```

run

Summary

Bayesian game-theoretic models provide a framework for:

1. Formalizing pragmatic theories (in this case Gricean pragmatics).
2. Making both qualitative and quantitative predictions about (possibly variable) language use and interpretation.
3. Capturing interactive co-construction of meaning (in this case truth-conditional).
   - The inference (15-b) arises as a product of coordination between the speaker and listener.

(15) a. Mary ate some of the cookies.
b. $\leadsto$ Mary didn't eat all of the cookies.

4. Capturing the contribution that speaker/listener prior beliefs make to pragmatic interpretation.

Let's explore. . .

Bayesian game-theoretic models provide a framework for:

1. Formalizing sociolinguistic theories (in this case Third Wave variation theory (Eckert, 2000, 2008, 2012)).
2. Making both qualitative and quantitative predictions about (possibly variable) language use and interpretation.
3. Capturing interactive co-construction of meaning (in this case social).
   - The inference (16-b) arises as a product of coordination between the speaker and listener.

(16) a. I have been work[in] on my paper.
     b. $\rightsquigarrow$ The speaker is friendly.

4. Capturing the contribution that speaker/listener prior beliefs make to social interpretation.

Game theory and sociolinguistic variation

A longstanding idea. . .

Tools/ideas from decision/game theory could be useful for analyzing language variation and change (Wittgenstein, 1953; Bourdieu, 1977; Dror et al., 2013; Clark, 2014, among others).

- Goffman develops these ideas (informally) (Goffman, 1961, 1967, 1970).

An epistemic perspective

Recent advances in cognitive science (i.e. Bayesian revolution) allow us to embed these ideas within a broader theory of linguistic and social cognition.

Matched Guise Technique

Matched Guise Technique (Lambert, 1967, et seq.)

An experimental method widely used in social psychology and variationist sociolinguistics to assess listeners' implicit attitudes towards speakers of different linguistic varieties.

- Participants listen to samples of recorded speech that have been designed to differ in specific and controlled ways.
- They one of two recordings (guises) which differ only in the alternation studied.
- After hearing a recording, participants' attitudes towards the recorded speaker are assessed (via interviews/focus groups and/or questionnaire/survey).

(ING)

Campbell-Kibler (2006, 2007): A MGT study with stimuli formed from the speech of 8 speakers.

Sample results

1. Speakers were rated as significantly more educated and more articulate in their *-ing* guises than in their *-in'* guises.
2. Speakers were significantly more likely to be described as a redneck in their *-in'* guises than in their *-ing* guises.

Social interpretation

Generalization from perception studies

Hearers make judgments about the properties that characterize speakers based on the linguistic forms that they use.

- Interpretation is only one side of the coin…

Generalization from production studies

Speakers strategically exploit hearer's interpretation process to construct personal linguistic styles, i.e. to communicate properties about themselves to their interlocutors.

Style Shifting

(Labov, 2012, 22) finds significant differences in President Obama's use of (ING) across three contexts.

- Casual: BBQing at a Father's Day BBQ on the White House lawn (72% -in').
- Careful: Answering political questions after the BBQ (33% -in').
- Formal: Scripted acceptance speech at the DNC (3% -in').

Obama's Use of (ING) (Labov, 2012)

Figure 3. President Obama's use of (ING) in three contextual styles.

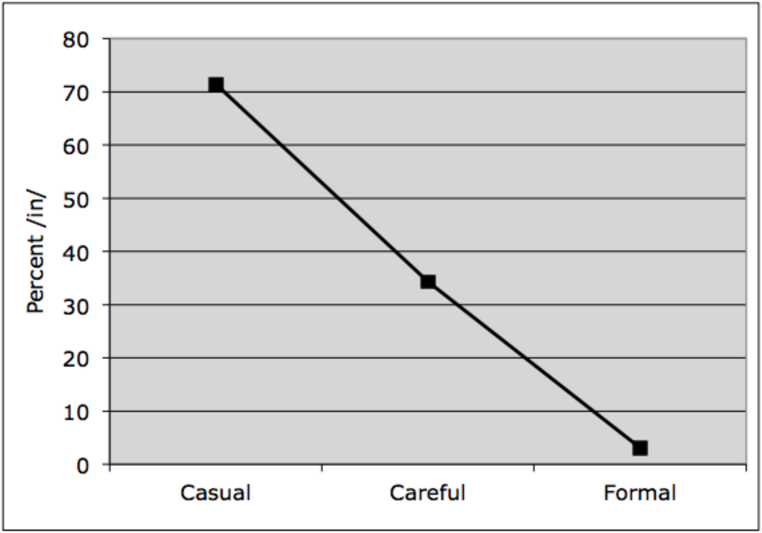

Convention

Why this pattern?

Labov (2012): As a community, we have conventionally associated meanings with *-in'* and *-ing*, which allow us to communicate extra information to each other through phonetic variation.

- *This consensus is publicly available and in one sense, understood by all. In the classroom, or on the pulpit, people will attribute the use of the -in' form to laziness, ignorance, or just plain rascality.*
- *Yet the high value we put on the -in' norm in other contexts is not hidden from public view. When we see the large illuminated sign,* DUNKIN' DONUTS*, we recognize the claim that dunkin' doughnuts taste better than dunking doughnuts...* (Labov, 2012, 22)

Sociolinguistic variation as rational language use

Summary

- Speakers assess how their speech will be evaluated by their interlocutors in a particular discourse context, i.e. the properties that they think their interlocutors will attribute to them. (interactivity)
- Speakers then choose the form that (they think) will be the most successful to construct their desired persona. (optimization/rationality)

Conclusion

Game theoretic tools are particularly well adapted to modelling this kind of linguistic communication.

Social Meaning Game (Burnett, 2016)

- S and L are the players.
- $\mathbb{P} = \{p_1, \ldots, p_n\}$ is a finite set of properties.
- $>$ is a relation on $\mathbb{P}$ that encodes antonymy.

Example: Obama across 3 contexts

$\mathbb{P}$ = {competent, incompetent, friendly, aloof}

(17) a. competent $>$ incompetent
b. friendly $>$ aloof

Personae

Third Wave Variation Theory focuses on how variants combine together (styles), which construct particular social types (personae) (see Podesva, 2004; Eckert, 2008; Zhang, 2008, among many others).

- Possible personae are collections of properties that *go together*.
- The personae are the set of largest consistent sets of properties.

Persona	Nickname
{competent, friendly}	'cool guy/gal'
{competent, aloof}	'stern leader'
{incompetent, friendly}	'doofus'
{incomptent, aloof}	'arrogant asshole'

Table: Universe in Obama example

Messages and Costs

1. $M = \{m_1, \ldots, m_n\}$ is the set of messages (i.e. variants) that S can pick from.
2. C is a function from M to the real numbers that assigns a cost to each message.
   - Way to incorporate linguistic/psychological constraints on variation...

(18) Today's Example

a. M = {ing, in'}

b. No cost difference between -ing and -in'.

Indexation and Indexical Fields

In Third Wave variation theory, individual variants have meaning that goes beyond their truth conditional meaning.

- Variants index sets of properties, called their indexical field (Eckert, 2008).

Variant	Eckert field
-ing	{competent, aloof}
-in’	{incompetent, friendly}

Eckert-Montague Fields

- In the spirit of Montague (1973), we can also look at indexical fields through the personae that they have the potential to construct.

Variant	Eckert field	Eckert-Montague field
-ing	{competent, aloof}	{comp., aloof}, {comp., friend.}, {incomp., aloof}
-in’	{incompetent, friendly}	{incomp., friend}, {comp., friend}, {incomp., aloof}

Table: Messages in Obama example

What does the speaker do?

S makes a hypothesis about L's beliefs concerning which persona(e) they instantiate: *Pr*.

- Suppose S thinks L doesn't have any opinion about them.

stern leader {comp, aloof}	cool guy/gal {comp, friend}	asshole {incomp, aloof}	doofus {incomp, friend}
0.25	0.25	0.25	0.25

Table: L has uniform priors.

Contribution of indexical fields

When they hear a variant, L restricts their attention to the personae in the (Eckert-Montague) fields.

m	stern leader {comp, aloof}	cool guy/gal {comp, frien}	asshole {incomp, aloof}	doofus {incomp, frien}
-ing	$0.\dot{3}$	$0.\dot{3}$	$0.\dot{3}$	0
-in'	0	$0.\dot{3}$	$0.\dot{3}$	$0.\dot{3}$

Table: L's beliefs immediately after hearing m ($Pr(P|m)$).

Speaker predictions: uniform priors

m	stern leader {comp, aloof}	cool guy/gal {comp, frien}	asshole {incomp, aloof}	doofus {incomp, frien}
-ing	1	0.5	0.5	0
-in'	0	0.5	0.5	1
Predict.	Cat. -ing	Var. (ING)	Var. (ING)	Cat. -in'

Table: S's predicted use of *m* given *P* (P_S(m\|P)).

Listener predictions: uniform priors

m	stern leader {comp, aloof}	cool guy/gal {comp, frien}	asshole {incomp, aloof}	doofus {incomp, frien}
-ing	$0.\dot{3}$	$0.\dot{3}$	$0.\dot{3}$	0
-in'	0	$0.\dot{3}$	$0.\dot{3}$	$0.\dot{3}$

Table: L's predicted interpretation of *P* given *m* ($P_L(P|m)$).

Listener predictions: uniform priors

m	stern leader {comp, aloof}	cool guy/gal {comp, frien}	asshole {incomp, aloof}	doofus {incomp, frien}
-ing	$0.\dot{3}$	$0.\dot{3}$	$0.\dot{3}$	0
-in’	0	$0.\dot{3}$	$0.\dot{3}$	$0.\dot{3}$

Table: L’s predicted interpretation of *P* given *m* (P_L(P|m)).

Do we ever really have uniform priors??

Obama at the BBQ

Suppose Obama wants to be perceived as the cool guy at the barbecue.

- He wants to construct the {competent, friendly} persona.
- He is worried about coming off as too aloof (since he is the president).

Obama at the BBQ

stern leader {comp, aloof}	cool guy/gal {comp, friend}	asshole {incomp, aloof}	doofus {incomp, friend}
0.30	0.20	0.30	0.20

Table: Obama worries about seeming aloof.

Obama at the BBQ (predictions)

Suppose $\alpha = 6$.

- P_{Obama}(-ing| {competent, friendly}) $\approx$ 0.31.
- P_{Obama}(-in'| {competent, friendly}) $\approx$ 0.69.

Obama after the BBQ

Suppose Obama is worried about coming off as incompetent when answering questions after the BBQ.

stern leader {comp, aloof}	cool guy/gal {comp, friend}	asshole {incomp, aloof}	doofus {incomp, friend}
0.20	0.20	0.30	0.30

Table: Obama worries about seeming incompetent.

Obama after the BBQ (predictions)

Suppose $\alpha = 6$.

- P_{Obama}(-ing| {competent, friendly}) $\approx$ 0.69.
- P_{Obama}(-in'| {competent, friendly}) $\approx$ 0.31.

Obama in front of Congress

Suppose Obama wants to be perceived as the stern leader in front of Congress.

- He wants to construct the {competent, aloof} persona.

Predictions

- P_{Obama}(-ing| {competent, aloof}) = 1.
- P_{Obama}(-in'| {competent, aloof}) $=$ 0.

Summary

- We can set proposals concerning social meaning, indexical fields and speaker/listener beliefs within formal Bayesian game-theoretic models.
- The models allow us to make qualitative and quantitative predictions for sociolinguistic variation and interpretation.

Obama at the barbecue is a toy.

- Listener prior beliefs and intended personae can be evaluated:
  1. Through questionnaires, in an experimental context.
  2. Through ethnographic analysis, in a sociolinguistic interview context.
- Our models can be extended with structure and dynamicity…

Evolutionary Game Theory

- Evolutionary Game Theory (EGT) models the propagation of some strategies over others across a population (Maynard Smith and Price, 1973).

Modelling the imperfective cycle

Requires developing a game for each synchronic stage individually and developing an overarching game that explains the transitions between these stages.

The former is like the Bayesian games we've already seen, and the latter requires a new addition: evolutionary dynamics.

The Progressive to Imperfective Path

- There are four stages in the cyclic diachronic process of innovating a progressive form and having it eventually generalize to the imperfective (Deo 2015's (24)):

X_{impf}	zero-PROG
$(Y_{prog})\ X_{impf}$	emergent-PROG
Y_{prog}, X_{impf}	categorical-PROG
Y_{impf}	generalized-PROG

- That means that there are three transitional states:
  - Emergence of a grammatical progressive (1 to 2)
  - Categoricalization of an optional progressive to an obligatory one (2 to 3)
  - Generalization of an obligatory progressive to more situations (3 to 4)

Synchronic Games

Speaker's Utility Function:

$$U_S(t,S,L) = \delta_t(S,L) - k \times n(s)$$

- The first part represents cooperation: it's 1 if L retrieves S's type and 0 otherwise
- *n(s)* represents the cost of having 2 similar expressions: it's 1 for 2 forms and 0 for 1 form
- *k* represents the value of how much S prefers communicative success over signal cost (low value for effective communication over signal cost)

Hearer's Utility Function:

$$U_L(t,S,L) = \delta_t(S,L)$$

Symmetry and Asymmetry

- So far we've seen asymmetric games in that the speaker and hearer play different strategies. But for the purposes of semantic change, it is important to consider each individual as both a hearer and a speaker.
- The way to turn this asymmetric strategy into a symmetric one is to calculate the utility for that person as a speaker and the utility for the same person as a listener and then combine them.
- The way in which Deo (2015) does this is to add both utility values together and divide by two.

Diachronic Games

- The evolutionary dynamics seek to model changes in the frequencies of different strategies in a population over time.
- S plays a particular strategy and is sequentially paired with other players in the population.
- The payoff obtained from each encounter is summed to yield the fitness of a strategy, and an average is taken that is weighted by the proportion of each type of participant that S plays against.
- This calculation of fitness is the rate at which players of that strategy are likely to replicate in the population, which can change the population composition over time.

Replicator-mutator dynamics

- In addition to the fitness calculation, we also need a measure of the learnability of individual strategies from the structure of the the learner's input.
- In Deo's model, mutations from one strategy to another happen because of input being misinterpreted during the acquisition process.
- Replicator-mutator dynamics are well-suited to processes where two competing strategies are at play (cf. Nowak et al. 2001)
- The replication rate of a strategy is the rate at which it might be adopted by players of other strategies.
- The mutation rate of a strategy is the set of barriers to the learnability of that strategy.

Conclusions

Game theory is a broad, flexible tool for analyzing events of strategic interaction.

- Signaling games allow us to incorporate the interactive aspect of communication into our formal pragmatic theories.
- Bayesian signaling games:
  1. Allow us to study pragmatic meaning and social meaning within a general theory of human cognition.
  2. Allow for the construction of mathematical models that make quantitative predictions for language use and interpretation.
- Evolutionary games:
  1. Allow us to model the changing behaviour of large populations over time.
  2. Provide a new perspective on grammaticalization.

Thanks!

Penny Eckert Devyani Sharma
Dan Lassiter

This research is funded in part by:

- The SSHRC Insight Grant *Variation et diglossie en français québécois: corpus, théories et modélisation* (PI: M. Tremblay).
- A visiting fellowship from the Center for the Study of Language and Information, Stanford University.
- The program “Investissements d’Avenir" overseen by the French National Research Agency, ANR-10-LABX-0083 (Labex EFL).

References I

Bergen, L., Levy, R., and Goodman, N. D. (2016). Pragmatic reasoning through semantic inference. *Semantics & Pragmatics*.

Bourdieu, P. (1977). The economics of linguistic exchanges. *Social science information*, 16(6):645–668.

Bourdieu, P. (1979). *La distinction: critique sociale du jugement*. Les éditions de minuit, Paris.

Burnett, H. (2016). Signalling games, sociolinguistic variation and the construction of style. *accepted in Linguistics & Philosophy*.

Campbell-Kibler, K. (2006). *Listener perceptions of sociolinguistic variables: The case of (ING)*. PhD thesis, Stanford University.

Campbell-Kibler, K. (2007). Accent,(ing), and the social logic of listener perceptions. *American speech*, 82(1):32–64.

Chater, N. and Manning, C. (2006). Probabilistic models of language processing and acquisition. *Trends in cognitive science*, pages 335–344.

Clark, R. L. (2014). *Meaningful games: Exploring language with game theory*. MIT Press.

Degen, J. (2015). Investigating the distribution of some (but not all) implicatures using corpora and web-based methods. *Semantics and Pragmatics*, 8(11):1–55.

Degen, J. and Tanenhaus, M. K. (2015). Processing scalar implicature: A constraint-based approach. *Cognitive science*, 39(4):667–710.

Degen, J., Tessler, M. H., and Goodman, N. D. (2015). Wonky worlds: Listeners revise world knowledge when utterances are odd. *Proceedings of CogSci37*.

Deo, A. (2015). The semantic and pragmatic underpinnings of grammaticalization paths: The progressive to imperfective shift. *Semantics and Pragmatics*, 8:1–52.

Dror, M., Granot, D., and Yaeger-Dror, M. (2013). Speech variation, utility, and game theory. *Language and Linguistics Compass*, 7(11):561–579.

Eckert, P. (2000). *Language variation as social practice: The linguistic construction of identity in Belten High*. Wiley-Blackwell.

Eckert, P. (2008). Variation and the indexical field. *Journal of sociolinguistics*, 12(4):453–476.

References II

Eckert, P. (2012). Three waves of variation study: The emergence of meaning in the study of sociolinguistic variation. *Annual review of Anthropology*, 41:87–100.

Frank, M. C. and Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998.

Franke, M. (2009). *Signal to act: Game theory in pragmatics*. PhD thesis, Institute for Logic, Language and Computation.

Franke, M. and Jäger, G. (2016). Probabilistic pragmatics, or why bayes' rule is probably important for pragmatics. *Zeitschrift fř Sprachwissenschaft*, 35:3–44.

Goffman, E. (1961). *Encounters: Two studies in the sociology of interaction.* Bobbs-Merrill.

Goffman, E. (1967). *Interaction ritual: essays on face-to-face interaction.* Aldine.

Goffman, E. (1970). *Strategic interaction*, volume 1. University of Pennsylvania Press.

Goodman, N. D. and Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in cognitive science*, 5(1):173–184.

Grice, P. (1975). Logic and conversation. *Syntax and Semantics*, 3:41–58.

Jäger, G. (2011). Game-theoretical pragmatics. In van Benthem, J. and ter Meulen, A., editors, *Handbook of Logic and Language*, pages 467–491. Elsevier, Amsterdam.

Labov, W. (2012). *Dialect diversity in America: The politics of language change*. University of Virginia Press.

Lambert, W. E. (1967). A social psychology of bilingualism. *Journal of social issues*, 23(2):91–109.

Lassiter, D. and Goodman, N. D. (2015). Adjectival vagueness in a bayesian model of interpretation. *Synthese*, pages 1–36.

Levinson, S. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT Press.

Lewis, D. (1969). *Convention*. Harvard UP, Cambridge.

Luce, R. D. (1959). On the possible psychophysical laws. *Psychological review*, 66(2):81.

Montague, R. (1973). The proper treatment of quantification in ordinary english. In *Approaches to natural language*, pages 221–242. Springer.

Ochs, E. (1992). Indexing gender. *Rethinking context: Language as an interactive phenomenon*, 11:335.

References III

Pinker, S. (1999). How the mind works. *Annals of the New York Academy of Sciences*, 882(1):119–127.

Podesva, R. (2004). On constructing social meaning with stop release bursts. In *Sociolinguistics Symposium*, volume 15.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423.

Shiffrin, R. M. and Steyvers, M. (1997). A model for recognition memory: Rem-retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2):145–166.

Silverstein, M. (1976). Shifters, linguistic categories, and cultural description. *Meaning in anthropology*, 1:1–55.

Sperber, D. and Wilson, D. (1986). *Relevance: Communication and cognition*. Blackwell, Oxford.

Steyvers, M., Griffiths, T., and Dennis, S. (2006). Probabilistic inference in human semantic memory. *Trends in cognitive science*, page 327?334.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT press.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022):1279–1285.

Van Rooy, R. (2003). Being polite is a handicap: Towards a game theoretical analysis of polite linguistic behavior. In *Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 45–58. ACM.

Wittgenstein, L. (1953). *Philosophical investigations. Philosophische Untersuchungen.* Macmillan.

Yuille, A. and Kersten, D. (2006). Vision as bayesian inference: analysis by synthesis? *Trends in cognitive science*, page 301?308.

Zhang, Q. (2008). Rhotacization and the ‘beijing smooth operator’: the social meaning of a linguistic variable. *Journal of Sociolinguistics*, 12(2):201–222.