Comments on the conference

Joel Sobel

Preface

These are a revised version of "notes to myself" prepared following the June 2023 Conference on Language and Game Theory.

What game theory doesn't do, I:

Daniel Rothschild pointed out that standard game-theoretic models of communication offer no connection between natural language and the meaning that messages acquire in equilibrium. I agree that this is a limitation of standard models,[1] but the game theory that comes from economics appears to deal with the problem in ways that are similar to the way that game theory used by linguists does. It assumes that an external language exists and makes some assumptions that connect the language to the communication game. This is the approach in Sidartha Gordon's presentation. His result parallels the RSA constructions for the kind of cheap-talk games that I understand. [Several papers by Andreas Blume parallels the approach of RSA.] There are different ways to integrate exogenous meaning into communication games (two examples: starting from an "honest" strategy and iterating best replies; making a restriction to monotonic strategies), but they exploit the same underlying structure.

My concern with RSA is that I haven't been able to use it to formulate general results. RSA techniques are not going to converge in many environments. When they converge, limits are likely to be sensitive to initial conditions (where there are multiple choices for sensible initial conditions).

I was surprised and happy that some audience members were comfortable with the monotonicity assumption that "does all the

[1] It is a problem shared by Farrell's idea too. He exploits the possibility that unused messages have default meanings, but does nothing to explain why the messages used in equilibrium take on particular meanings.

work" in the results Sid described.[2] Game theorists are more skeptical. And the skepticism is justified for a "linguistic" and a mathematical reason. The linguistic reason is that the monotonicity assumption applies only in restrictive environments (the domain of uncertainty is one dimensional and preferences satisfy a sorting condition). The mathematical reason is that all assumptions are bad.

Just as I'd like to see general statements about RSA limits, I'd like to be able to talk about more general restrictions on strategies that capture the spirit of monotonicity.

Daniel's criticism, interpreted narrowly, is not justified. (That is, there are attempts to incorporate some notion of natural language into the game-theoretic analysis of communication. Indeed, doing so is essential if one wishes to talk about some issues, like lying, that are critically important to economists.) Daniel's criticism, interpreted broadly, is justified. It is related to a criticism he raised on Friday.

What game theory doesn't do, II:

Daniel mentioned another limitation on Friday in response to Herakles Polemarchakis's question on Friday. He pointed out that "messages" in game-theoretic models have little relationship to real language. With the exception of a few articles (including work by Andreas), game-theoretic messages are tokens that are statistically related to unknown information. They acquire meaning through the magic of rationality and equilibrium assumptions. They lack even the simplest properties that a normal person would expect from language. (I pointed out that they don't distinguish reference to states from orders to take particular actions, which you could hope to incorporate in a formulation that didn't include grammar.) Game theory so far has nothing to say about the structure of language. It is somewhat surprising that the current models can say anything interesting to students of language.

There are fascinating and deep questions about how "tokens" can be combined to make more complicated utterances. One can supplement observation by mathematical modeling to formulate questions

[2]My delight was moderated after I learned that assumption that there is an upper bound to the highest message would be unacceptable to linguists. Ultimately, I rationalized this limitation because (a) without the assumption, the conclusions of the analysis don't hold; and (b) conclusions break down for a "natural" reason: people will keep trying to use higher and higher superlative. I'd welcome a reference to something that discusses (what we call) monotonicity restrictions.

and perhaps obtain answers, but – at least to a first-order approximation – these models may combine elements of logic, learning theory, and computational complexity.

There is no reason to expect game theory to help on these questions. The issues are fundamental. Consequently, I can understand people in the linguistics community rejecting game theory as a tool. But I think that language and communication are essential for social interaction, that social interaction is inherently strategy, and that therefore it offers opportunities for strategic interaction.

What game theory doesn't do, III:

Heather Burnett raised concerns on Friday that indicate a different limitation. Heather knows that many of the kinds of interactions that she wishes to study are interactive decision problems. She knows in principle that game theory is a technical tool for studying this kind of problem. But she hasn't seen anything from game theory that helps her understand what she wants to understand. She has strong hypotheses on how you speak influences how people view you. I could imagine that a game-theoretic model might her organize some observations and make predictions, but I suspect that while doing that might make it easier for me to understand what she knows, I'm not sure it would help her know more. She has an adequate conceptual framework and wants to collect data. It would be exciting to provide a theoretical structure that organized what she observes and makes novel predictions. It is less exciting to simply engineer a model that conforms to observation.

What I learned, I: Vagueness

I enjoyed Paul Egré and Benjamin Spector's paper.[3]

I have not settled on a definition of vagueness. Instead I'll discuss properties that I find interesting. I do have an operational definition of precision: this is a situation in which Sender's message has a clear interpretation – it induces a non-degenerate distribution over states: Associated with m, there is a corresponding belief $\mu(\cdot \mid m)$ defined over states. The trouble with this definition is that every

[3]I studied Bart Lipman's paper years ago and found it stimulating. In some sense, the only connection between the papers is that they use the same term, but they define it differently and do different things.

message in a game-theoretic equilibrium will be either fully revealing or imprecise. There is no room for vagueness or ambiguity. So what is vagueness about? One possibility is that it is a message that means different things in different contexts. For this one needs a natural set of messages (exogenous language), some parameter that describes "context", and comparable state spaces. For example, m can be the word "small," θ can represent height (that is, a positive real number), and context could be whether I'm talking about basketball players, three year olds, or giraffes. Benjamin mentioned this interpretation, but I gather that his analysis doesn't directly apply to it.

Another possibility is that a statement is vague if different people interpret the statement differently. (But maybe it would be better to call this ambiguity.) My first stab at how to model this is to assume that the audience has private information. When you combine the message with private information, you get a different interpretation. Formally, this approach seems close to the previous one – "context" and "private information" both serve as a hidden parameter that influences beliefs.

Paul and Benjamin provide a way to interpret vague messages. That is, they formally translate "about" into a statement about beliefs. I don't think the result (that vague messages may be more effective than imprecise ones) depends on the exact way in which "about" is translated into a belief.[4] But "about" is always translated into a specific probability distribution. I'm inclined to call it imprecise. Further, if the speaker has a particular distribution in mind, why doesn't she just report it. When we use a vague term is it just a simple form of a report of a belief?

One reason for vague statements is that the speaker herself lacks hard information. I cannot honestly report that there are between 10 and 20 people in the room. I'm not sure. Another is that vague words are flexible (they apply in different contexts). I don't know how to formalize it, but surely there is an advantage to be able to make relative statements that make sense in a large number of domains. Another is that the listener only needs vague information and vague statements are easier to interpret. If you only need to know what day it is, I'm doing you a favor (maybe a very small

[4]And this is fortunate. If someone told by that a person was "about" 90 years old, my point estimate would be lower than ninety. If ninety doesn't do it for you, what about 110? Similarly if someone was about 2 meters tall.

favor) if I say "after midnight" rather than 12:04 because the first statement tells you exactly what you need to know and the other contains too much information. But "after midnight" can be construed as imprecise but not vague (although to me "after midnight" suggests somewhat close to midnight rather than any time after midnight). These concerns seem orthogonal to the driving force for efficient vagueness in Paul and Benjamin's paper.

Suppose that two couples are arriving at the airport. I know everyone's height and the identity of the person you should contact. I say: "talk to the man who is taller than his wife" instead of "talk to the man who is 5 feet 9 inches tall." I do this because I know that you can make comparisons but not identify absolute heights. Is this vague?

What I'd like to see is a model that operationalizes the idea that some messages are easier to encode (going from belief about state to message) or decode (going from message to belief about state) than others.

Aside: Describing conditions under which vagueness is efficient is one way to justify the existence of vague communication. It is not the only way. You could image that language use is the result of an evolutionary process or the equilibrium of a non-cooperative game. In both cases there is no guarantee that outcomes would be efficient.

What I learned, II: Rational Speech Act Theory

I heard about RSA prior to the conference and knew about Michael Franke's early work and the Frank-Goodman announcement in *Science*. The ideas behind it are similar to techniques used by "my people." The basic idea is to identify the prediction of a communication game with the limit of best reply dynamic (or noisy best-reply dynamic, which I guess is called soft max) starting from the initial condition of "truth." For games with structure (supermodular games and their generalizations), this procedure is closely linked to iterative deletion of dominated strategies.

The psychology/linguistics literature does not seem to worry about proving convergence. It looks at richer examples than the typical economics approach, computes limits empirically, and has some success stories (the limit is descriptive). This procedure is quite similar to what Andreas did in his talk (and other papers) and what Sid

(and co-authors) did. The economists in game theory are (in my view) more careful about technical details. The linguists have sexier examples. Because of these different focuses, each group has something to offer the other.

What game theory might do, introduction

Models from games and decisions won't be useful to most linguists. The issue is whether the models are useful to some linguists. More specifically, is there a reason to believe that people trained in linguistics who are interested in strategic interaction and comfortable with game theoretic techniques have something to gain from talking more to game theorists who work in other disciplines (from the US perspective, this means economists) and whether game theorists who study communication but have no expertise in language have something to gain from talking to linguists? I am trying to make the case that the answer is yes. The answer may be no: Linguists interested in formal models of games and decisions may know the techniques well enough so that they don't need "help" from technical people in another discipline. Economic game theorists have their own problems and may not be flexible enough to be able to identify central problems for another discipline. I am convinced that there are some clear gains from having more interaction. From my perspective, some of the applications of games to language have "reinvented the wheel" and the rediscoveries have been inferior to the original. Personally (and possibly professionally), I have benefited from trying to come to terms with Austin and Grice.

What game theory might do, repeating my talk

In my talk, I described a few topics that I find interesting that may be on the border of language and economic theory. The three "organizing principles" were: it is essential to take into account "complexity;" it is not essential to focus on conflict of interest; there are interesting questions in which "language" is unnatural. The third principle means that I am willing to sacrifice most (all?) questions that linguists might care about. So maybe my interest in "codes" has little overlap with topics that would interest all of us. (My justification is that understanding how the way we communicate changes when we take into account that what we talk

about changes is a topic of broad interest and studying artificial language is a way to make the problem tractable.) Figuring out how to model "complexity" is a problem for economic theory (because classical economic theory assumes that people can optimize without cost). Behavioral economics has pushed the profession towards taking bounded rationality seriously, but there are not yet general techniques, so I am motivated by the general question of how to model complexity. The questions motivated by (my view of) language serve as examples.

My Questions

I raised some questions in my talk (thinking about Grice with explicit conflict; differentiating between reference and connotation; understanding the nature of convex categories; the effect of changing environments on the complexity of language). After the conference, I had a renewed interest in defining and understanding vagueness and ambiguity, a renewed awareness that there are often systematically different interpretations of statements that have the same semantic content; and (based on no specific presentation) curiosity about whether it would be useful to have a strategic model of the wide variations of informal/formal distinctions (tu/vous) in natural language.

Coda

I view applied game theory (and, when I am being grandiose, all theory) as a way for uninformed people to appear to be smart. I have not systematically studied the field. I only speak one natural language (and do so poorly). I have no mastery of language use in any particular setting. But I can manipulate formal models. Formal models permit one to make qualitative statements about many, seemingly unrelated, situations. A good model can organize a lot of seemingly different bodies of evidence. It allows you to make predictions about novel environments. So someone with a good model can sometimes make intelligent statements about situations he knows little about. Someone with a good model can absorb new information more quickly than someone who has more specialized information but lacks an organizing structure. Models, of course, are abstractions. A model-bound observer will neglect information

that does not fit the model and will strain to force situations into a rigid framework. Bad things happen when theorists forget this.