Language and Economics[1]

Barton L. Lipman

Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706

email: blipman@ssc.wisc.edu

May 2002

[1] I thank Andreas Blume for helpful references and Eddie Dekel and Duane Seppi for numerous discussions which influenced the views expressed here. I apologize to Ariel Rubinstein for using (a permutation of) the title of his book for the title of this essay.

1 Introduction

I find myself in the rather awkward position of trying to survey a literature which (as I will feel compelled to define it) almost doesn't exist. Outside of Rubinstein's [2001] *Economics and Language* and a handful of other papers, economists seem to have largely ignored language.

This is an unfortunate state of affairs. The world people live in is a world of words, not functions, and many real phenomena might be more easily analyzed if we take this into account. For example, consider incomplete contracts. Our models treat contracts as mathematical functions and hence find it difficult to explain why agents might not fully specify the function. Of course, real contracts are written in a language and may not unambiguously define such a function — not its domain or range, much less the function itself. More broadly, language is the means by which we communicate with others (or with the self) for information, coordination, persuasion, etc. Given that we interact in language, we should expect language to affect and be affected by our interactions.

To be fair, my claim that this topic has been neglected rests in part on how I choose to define the topic. If "language" is taken to refer to any form of information transmission, there is a huge literature on the subject. However, while this literature has generated many fascinating and important insights, it has had relatively little to say about language *per se*. To clarify the point, consider Spence's famous 1974 model of labor market signaling. In that model, education costs which are negatively correlated with productivity make the level of education a credible signal of ability. Is this a model of language? In a sense, yes: education is the medium by which information is communicated. On the other hand, it says little about spoken or written words. To the extent that it is, at least in part, the medium we are interested in, a model where communication takes place through the medium of costly actions says little about language.

Similarly, the literature on mechanism design and implementation involves many "message sending" games. However, as with Spence, the focus tends to be on outcomes, not the language by which information is communicated. The same could be said of other seminal papers such as Crawford and Sobel [1982] or Farrell [1993]. In all cases, language is highly relevant but not really the focus of the research.

In this essay, I will even ignore a body of literature which does discuss language explicitly: the game–theoretic literature on how language acquires meaning. (See, *e.g.*,, Blume, Kim, and Sobel [1993] and Wärneryd [1993].) Essentially, these papers consider how evolution can lead agents to solve a rather complex coordination game regarding which vocal sounds become associated with which concepts. While an interesting issue, this doesn't say much about what sort of coordination will emerge. That is, there are no results (that I know of) regarding the likely nature of the language which emerges from such a process.

The papers I do discuss are those which seem to me primarily focused on the medium of communication. Since it is highly subjective whether a given paper focuses "enough" on the medium to fit my purposes, it is perhaps more accurate to say that I will define the relevant literature to be the set of papers I discuss.

More specifically, I will discuss four topics. In Section 2, I discuss work on the "optimal" structure of language. Specifically, I discuss Rubinstein's work on binary orderings in natural language, Chapter 1 of his book, and Blume [1999, 2000]. In Section 3, I discuss effects of language on choice. In particular, I discuss Lipman [1999] (though I reach a negative conclusion here) and Chapter 4 of Rubinstein's book. In Section 4, I discuss work on debate and inference rules in this context. There are more papers on this topic and I will primarily discuss Lipman–Seppi [1995], Glazer–Rubinstein [2001] (summarized in Chapter 3 of Rubinstein's book), and Spiegler [2002b]. Finally, in Section 5, I discuss vagueness, a topic which has received virtually no attention but seems to me to be of great importance.

2 Optimal Language

While this perspective may seem odd to some who study language, an economist's first instinct in thinking about why language has the structure it does is to consider what it optimizes.[1] I know of three intriguing papers which take this perspective, one in Rubinstein's book and two by Andreas Blume.

Altogether, these papers suggest four interrelated criteria which may be part of defining the optimality of a language. I describe these informally and then turn to more formal meanings below.

One criteria, called *indication–friendliness* by Rubinstein, is the ability of the language to identify objects. Clearly, if each object in the universe has its own name, then an agent can unambiguously refer to an object by its name. In reality, we explain which of a set of objects we are referring to by giving some partial description of it. If a language is too limited, it may be impossible to unambiguously identify some objects. A language is indication–friendly if it does not suffer from this problem.

A conceptually similar idea is how well the language indicates relationships among the objects being discussed. Rubinstein refers to this as *informativeness*.

[1] I will not pretend to know the linguistics literature. A very casual reading suggests that optimality considerations do play a role in some of the research in the area.

A different issue is the ease by which the language is acquired. There are many versions of this one could consider. Rubinstein and Blume consider essentially the same idea, termed *describability* by Rubinstein. The idea, as Rubinstein sees it, is to postulate that language is taught by a speaker to a nonspeaker. In this case, the teacher will choose which parts of the language to communicate. A language is more readily describable if there is less that the teacher must say to convey the language to the pupil.

Finally, Blume discusses the *creativity* of a language. By this, he means the ability of the language to communicate novel and unexpected events. Clearly, real languages have the flexibility to describe situations that were never recognized as possibilities while the language was evolving. Languages which are more flexible in some sense are more able to do this.

To move toward more formal definitions, let Ω be a set of objects, the objects which are to be discussed in the language. Rubinstein and Blume differ in the aspect of language they focus on. Rubinstein considers the role of binary relations in language, so he takes a language to be a binary relation, say R, on Ω. I discuss Blume's formalism below.

Rubinstein says that a language is indication–friendly if for every set of objects $A \subseteq \Omega$ and any object a in the set, there is a way to identify a from A. Somewhat more formally, there must be a formula in predicate calculus with one binary relation and no constants which a satisfies and which no other element of A satisfies. The phrase "no constants" rules out the possibility that the names of the objects are itself part of the predicate calculus, a possibility which would, of course, trivialize the result. Rubinstein shows that a language R is indication–friendly if and only if R is a linear order — that is, if and only if R is complete, transitive, and asymmetric.

It is not hard to see that R is indication–friendly if it is a linear order. For any set A and element $a \in A$, a has a unique rank within A according to the order R. This can be stated as a formula in predicate calculus using R and hence can be used to identify a from A.

To see that it is necessary that R be a linear ordering, suppose R is indication–friendly. Then from the set $\{a, b\}$, there must be a formula which selects a and another which selects b. If $a \ R \ B$ and $b \ R \ a$, then any formula involving R and A which doesn't include the names of the objects is true about a if and only if it is true about b. Hence we cannot identify either object from this set, contradicting indication–friendliness. Hence R must be asymmetric. If we have $a \ \not R \ b$ and $b \ \not R \ a$, we have the same problem. Hence R must be complete. Finally, consider the set $\{a, b, c\}$. Suppose $a \ R \ b$, $b \ R \ c$, but $a \ \not R \ c$. By completeness, $c \ R \ a$, so we have a cycle. But then any statement involving only R and A which is true about one element must be true about all since any element is a permutation away from any other.

Informativeness is conceptually similar to indication–friendliness in that both refer to how well the language can identify the things being discussed. The difference is that indication–friendliness concerns identifying individual objects in a set, while informativeness concerns identifying relationships among the objects. Suppose that we must choose a relation R which will be the language before learning which binary relation will be the one giving the relationships we actually need to describe. If the relation we actually need is $R' \neq R$, then use of R will imply some mistakes in describing R'. What R minimizes the probability of inaccurate description?

To state the question more precisely, think of a relation as a subset of $\Omega \times \Omega$ minus the diagonal giving the set of (a, b) such that $a\ R\ b$. Rubinstein considers only asymmetric orderings, so we remove pairs of the form (a, a) throughout. Given any other relation, say R', we have four ways to use R to describe R'. First, there are two trivial ways: to say we always have $a\ R'\ b$ and to say we never have $a\ R'\ b$. Second, we can describe R' as if it were R or as if it were the negation of R. That is, we can say that $a\ R'\ b$ iff $a\ R\ b$ or iff $a\ \not R\ b$. Rubinstein assumes that when the needed relation is determined, we will describe it using whichever of these four approaches leads to the fewest errors. In other words, we will choose from $\{\Omega \times \Omega, \emptyset, R, R^c\}$ to minimize the number of ordered pairs outside the intersection of the chosen set and R' (where R^c denotes the complement of R in $\Omega \times \Omega$ minus the diagonal). Given this, we choose R to minimize the expected number of errors given a uniform prior over R'. Rubinstein also wants R to do well given some restriction of the set of objects Ω.

Rubinstein shows that it is approximately optimal to choose a binary relation which is complete and asymmetric. To see the intuition, suppose R is complete and asymmetric. Then it splits $\Omega \times \Omega$ minus the diagonal into two equal sized pieces. That is, for any pair a, b with $a \neq b$, such a relation says either $a\ R\ b$ or $b\ R\ a$ but not both. Hence one order of this pair is placed in R and one in R^c. Thus R and R^c have the same number of elements. Furthermore, this would be true for any subset A of Ω. Given the symmetry of the problem, it should seem plausible that a symmetric splitting is at least approximately optimal, as Rubinstein shows.

Next, I turn to describability. As formulated by Rubinstein and Blume, the idea is to consider the minimum number of examples needed to learn a language. In Rubinstein's model, a language is a binary relation, so this means the minimum number of pairs we need to know the comparison of in order to deduce the entire relation. Of course, if one knows nothing of the structure of the relation, one has to learn all comparisons, making all binary relations equally difficult to learn. So Rubinstein measures describability of a relation by looking for the minimum number of examples needed given some description of the relation in predicate calculus. The idea is that a relation which requires relatively few examples is one easily taught and hence more likely to be passed on.

Note that if we are told that the binary relation we are supposed to learn is a linear ordering, then $\#\Omega - 1$ examples is sufficient to learn it (where $\#$ denotes cardinality). That is, we know the entire relation if we are told the comparison between the top ranked object and the second, the second and the third, etc. Rubinstein shows that if we restrict attention to complete, asymmetric orderings, then linear orderings are approximately optimal for large Ω sets. More specifically, if Ω is large, then the minimum number of examples for any complete, asymmetric ordering must be approximately $\#\Omega - 1$.

In Blume [1999], a language is a function from Ω to some set of "words" M which associates a different word with each different object. That is, it is a function $f : \Omega \to M$ such that $\omega \neq \omega'$ implies $f(\omega) \neq f(\omega')$. As in Rubinstein's model, Blume assumes a certain structure to the information the learner has regarding the language to be learned. He models this by supposing that there is a set of possible languages, $\mathcal{F}$, and that the learner initially knows only that $f \in \mathcal{F}$. Learnability is defined in terms of the set $\mathcal{F}$ and the ability to identify the correct f from this set. Blume calls a pair of the form $(\omega, f(\omega))$ an observation and defines the learnability of a language to be the minimum number of observations needed to unambiguously identify f within $\mathcal{F}$. More specifically, if $A \subseteq \Omega$ is a set of objects, A identifies $\mathcal{F}$ if for every $f, f' \in \mathcal{F}$ with $f' \neq f$, there is some $\omega \in A$ such that $f(\omega) \neq f'(\omega)$. That is, if the agent knows the appropriate words for the objects in A, this is sufficient to know the language. Learnability, then, is minimum cardinality for a set which identifies $\mathcal{F}$. Clearly, this is conceptually the same idea as Rubinstein's describability.

Blume puts more structure on the problem and uses it to obtain a different kind of result. He assume that Ω is a product space and that M is a set of strings of symbols. To illustrate, assume that $\Omega = \{a, b, c\} \times \{x, y, z\}$ and that the symbols available are all the Greek letters. That is, M consists of the set of all finite sequences of Greek letters. One language, then, would be $f(a, x) = \alpha$, $f(b, x) = \beta$, $f(c, x) = \delta$, $f(a, y) = \gamma$, etc. In other words, we have one letter for each pair in Ω. Another language would be $f'(a, x) = (\alpha, \xi)$, $f(b, x) = (\beta, \xi)$, $f(a, y) = (\alpha, \varphi)$, etc. In other words, we have one letter for each of the components a, b, c, x, y, and z and describe a pair by the pair of letters they correspond to.

Suppose the agent knows only that the language is of the former type where there is one letter for each pair. In other words, $\mathcal{F}$ is the set of languages mapping Ω into the Greek letters. Since there are nine points in Ω, it will take nine observations to identify $\mathcal{F}$. Suppose by constrast that the agent knows only that the language is of the latter type. Now he needs only three observations to learn the language. More specifically, if he observes the names of the objects (a, x), (b, y), and (c, z), he will identify the Greek letters being assigned to each of a, b, c, x, y, and z and hence will know the language. In this sense, a language which respects the product structure — is *modular* in Blume's terminology — is easier to learn than one without this structure. He shows that, in a certain sense, such modular languages are the easiest to learn.[2]

This is related to Blume's [2000] notion of the creativity of a language. To see the idea, let us continue with the example above. First, suppose that the language is the one which identifies each of the nine points in Ω with a different Greek letter. Then, as noted, one has to observe *all* objects and their names to know the language. If one produces an object never seen before, the language learner cannot possibly know its name.

By contrast, consider the modular language. Suppose the learner has only observed the names given to (a, x) and (b, y). Then even though he has never seen (a, y), he knows its name since he now knows the letters associated with both a and y. In this sense, the language is capable of communicating about novel observations, ones which have never been seen before. Unfortunately, I cannot provide a brief and semi–comprehensible summary of the results, so I encourage the interested reader to consult Blume's paper.

3 Language and Actions

If language is the medium of thought, then the nature of language may well have an effect on the way people choose. There are at least[3] two distinct forms such an effect could take. First, language could *direct* decision making. Second, language could *constrain* decision making. By the former, I mean that language can push the agent's deliberations in a particular direction. By the latter, I mean that decisions which are made in language presumably must be formulated in language, a requirement which can constrain decisions.

As to the former, there is much evidence but little theory. The experimental work on framing effects clearly shows that the phrasing of the choice problem can affect the decision in a nontrivial fashion. For example, Shafir, Simonson, and Tversky [1993] consider choices between two options and demonstrate that the decision is strongly affected by whether the question is which option to accept versus which option to reject. Their argument is that the former phrasing leads subjects to look for reasons why they should like one of the options, while the latter leads them to focus on why they should dislike one of the options.

I know of no model of this effect. At one point, I had hopes that the model in Lipman [1999] might shed light on framing, but am no longer optimistic. In that paper, I considered an agent who receives information in the form of statements in some logical

[2]Blume [2000] includes some more general results on learnability of languages.

[3]Spiegler [forthcoming, 2002a] discusses another way language may affect choice. He models players in a game choosing strategies which they can justify subsequently in a "debate." I do not discuss these papers here because the model does not include language *per se*.

language, but may fail to recognize logical equivalence. More precisely, there is a set of possible information sets (collections of statements) the agent can receive, $\mathcal{I}$. Given any $I \in \mathcal{I}$, the agent has a preference, $\succ_I$, over a set of acts. I assumed an objectively given logical structure relating these information sets to the consequences of the acts. (The details are irrelevant here.) The key idea in the model was that the agent's preference in response to information I could differ from his preference in response to information I' even though these are logically equivalent. This seems like exactly what framing is about: differences in phrasing, not content, affect choices.

The approach I took was to model this perception as an information effect. While arguably useful for some purposes, I don't think this approach does much to improve our understanding of framing. To see the point, consider the classic experiment of Tversky and Kahneman [1981] who found that doctors responded more cautiously to information about alternative treatments phrased in terms of death rates than when this same information was presented in terms of survival rates. Suppose the two treatments are A and B. Let d denote the data given the doctors in terms of death rates and s the same data phrased in terms of survival rates. For concreteness, suppose data d leads the doctor to choose A, while s leads him to choose B. That is, $A \succ_d B$ and $B \succ_s A$. My approach was to model this by supposing that in the agent's mind, d and s are not equivalent. That is, there is a situation the agent conceives of in which, say, d is true but s is false — for example, perhaps d is true in a state in which a death leads to a lawsuit while s is not true there. When the agent is told that s is true, he rules out this state as possible but does not do this if told d is true.

While we can model the choice this way, it isn't clear that it is fruitful to do so. The idea of framing is that different statements call to mind different possibilities or lead to different ways of considering the problem. While these differences might be modeled as information, this seems to tell us little about why these differences arise or what kinds of differences are likely to be important. For example, this approach does not give an obvious way to think about the Shafir–Simonson–Tversky example above where the presentation of the data is held fixed but the way the question is phrased varies.

There has been more success in modeling how language may constrain decision making. In Chapter 4 of *Economics and Language*, Rubinstein gives two intriguing examples illustrating how this may work. Intuitively, if people do think in language, then it seems plausible to say that their preferences or decision rules should be expressible in this language. Even if they do not think in language, it may be necessary to formulate a decision rule in words in order to convey it to an agent who will carry it out. I will discuss the implications of this definability requirement in one of Rubinstein's examples. I don't wish to go into too much detail, so the reader is warned that my statements of his results are a little imprecise.

In the example, Rubinstein considers the preferences of an agent in a world of multi–attribute decisions. Specifically, each option available to the agent has a vector of attributes. Rubinstein assumes that the decision maker has clear preferences over each individual attribute, but must aggregate these to generate his preference over options. Rubinstein asks what kind of preferences can emerge if the agent must be able to state his aggregation rule as a formula in propositional calculus with the attribute preferences as atomic formulae. Rubinstein shows that if we require the aggregation procedure to always generate transitive and complete preferences, then the procedure must be lexicographic. That is, there is an ordering of the attributes such that a lexicographic comparison of the attribute vectors generates the overall preference. While the simplicity of lexicographic comparisons is clear, the uniqueness result is surprising, to say the least.

Putting together the concepts of this section and the last, we find an intriguing circularity. In the previous section, I discussed papers based on the idea we structure language in a way which seems useful to us given our perception of the world. On the other hand, the papers discussed in this section hypothesize that, once developed, language affects the way we see the world. An analysis which includes feedback effects could be quite interesting.

4 Debate

In the previous two sections, I considered the structure of language and the connection between language and actions. So far, I have said little about the meaning of language or how people interpret statements. This turns out to be particularly interesting in the context of a debate with evidence. Evidence gives statements a "pure" or intrinsic meaning which can interact with interpreted meaning in surprising ways.

To illustrate this point and more broadly the power of language, consider the following simple model, a slightly modified version of Lipman–Seppi [1995]. There is a finite set of states of the world, S. Each of two debators, 1 and 2, know the true state. The observer does not. The debators present evidence to the observer, who then chooses an action $a \in A$ which affects the payoffs of all of them. The observer's optimal action depends on the state. To make the results particularly stark, I assume that the preferences of the debators regarding the observer's action are independent of the state so, in general, they have no incentive to be honest.

I assume, however, that they can present evidence which proves some of their claims. Formally, the message set available to debator i in state s is $M_i(s)$. To understand this, note that presenting a message feasible only in state s proves that the state is s. In this sense, such a message constitutes proof. More intuitively, if a house deed is presented with a particular individual's name on it, this proves the person owns or has owned a home.[4]

I begin with an example where only the second debator has any evidence. More specifically, I assume that $M_1(s) = S$ for all $s \in S$. That is, debator 1 can name any state of the world as his "claim" of what is true, but has no evidence whatsoever to produce. I will assume that debator 2 cannot conclusively prove which state is the true one, but can rule out any false state. That is, for every s, there is a message $m_{\neg s}$ such that $m_{\neg s} \in M_2(s')$ if and only if $s' \neq s$. Hence $m_{\neg s}$ proves that the state is not state s and nothing more.

For concreteness, assume $S = \{1, \ldots, 1000\}$, that the observer's action set A is equal to S, and that the optimal action in state s is s. In this sense, we can think of the observer as trying to guess what the state of the world is.[5] Assume that the payoff of debator 1 is strictly increasing in the action chosen by the observer, while 2's payoff is strictly decreasing. For example, the observer may be a judge who must determine the amount of damage done (and hence compensation owed) by 2 to 1.

Suppose there is time to only hear two messages, both from 1, both from 2, or one from each in some order. It is tempting to conjecture that debator 1 cannot possibly influence the observer's action. After all, debator 1 has no evidence and no apparent incentive to tell the truth. This intuition suggests that the optimal choice by the observer is to request two messages from player 2 and none from player 1. However, this choice gives the observer very little information. It is not hard to show that in any equilibrium when only debator 2 speaks, there must be (at least) one pair of messages for which the observer's inference is only their pure information content. Since any pair of messages is feasible in at least 998 states, this means that we must have at least this many states pooled in any equilibrium.

On the other hand, if the observer does solicit a message from 1, there is a separating equilibrium! To see how this works, suppose 1 sends a message first, followed by a reply from 2. Let 1's message be interpreted as a claim of what the true state of the world is. If 2's evidence does not refute 1's claim, then the observer concludes that 1 honestly revealed the state. However, if 2's evidence does refute the state 1 claimed, the observer draws the worst possible inference for 1. (Alternatively, we can allow debator 2 to also make a claim of a state which the observer should infer in this situation.) It is easy to see that 1's optimal strategy given this inference rule is to honestly reveal the state. 2

[4]This statement assumes, of course, that forgery is impossible. If forgery is possible, the presentation of the deed proves that either the person has owned a home or someone has committed forgery. In this case, the deed constitutes less definitive evidence, but its presentation does rule out some states of the world.

[5]One can easily modify the example to make it into a debate between two competing positions — say whether the true state is in $\{1, \ldots, 500\}$ or $\{501, \ldots, 1000\}$. I omit the details for brevity.

cannot refute a true claim, so the observer will infer the state correctly.

Startlingly, then, 1's evidence–less messages are worth more than a great deal of evidence. More precisely, suppose the observer can choose between seeing one message from 1 followed by one message from 2 versus, say, 500 messages from 2. The argument above shows that the observer is better off with the former. In this sense, it is the combination of words with evidence that is critical.

Lipman–Seppi [1995] show how inference rules analogous to the one used here can generate full information revelation in a surprisingly wide set of circumstances. Our primary focus is on *BUR rules* where BUR stands for "believe unless refuted." More precisely, assume we have many debators who speak in some order. The exact order can be fixed or can depend on the statements made. For example, it could be that 1 speaks first and then either 2 or 3 speaks next, depending on what 1 says, followed by the other. We require that each debator get at least one chance to speak regardless of what anyone says.

The observer need not know the preferences of the debators. We require only that the debators have *conflicting preferences*. More specifically, we assume that given any two possible conclusions the observer could reach, we can find one debator who prefers the observer to reach the first conclusion and one who prefers the second. More specifically, if a and a' are actions the observer might take for some beliefs, then we can find one debator who strictly prefers a and one who strictly prefers a. To simplify the discussion, I will phrase everything in terms of the state inferred by the observer without specific reference to actions. The observer knows nothing about the preferences of the debators except that they satisfy this property.

For simplicity, I will assume that every debator has access to the same evidence. That is, $M_i(s) = M(s)$ for all i. I also assume that every debator sends a message consisting of two components: an element of S and a piece of evidence from $M(s)$. The first component is interpreted as the state the debator claims is the true one.

The BUR rule is the following. If the first debator's claim of a state is accompanied by evidence satisfying a certain burden of proof, then this claim is conditionally accepted. It remains conditionally accepted as long as no debator subsequently refutes it. When some debator does refute it, his claim, again if accompanied by appropriate evidence, is conditionally accepted. We continue in this fashion through the rounds of debate. The last conditionally accepted claim is the observer's inference.

When claims are not accompanied by evidence satisfying the burden of proof, the observer simply interprets the claim as different from the one stated. Given the assumptions I make below, it is always possible to find some state consistent with the evidence presented so far for which this debator's evidence does satisfy the appropriate burden of proof. The exact way the observer chooses among these alternative interpretations is irrelevant for our result so long as it is known to the debators. Given this, I will ignore the possibility of such deviations in the discussion here.

To explain the burden of proof, first consider two states with $M(s) \subset M(s')$ (where $\subset$ means strict inclusion). If the true state is s, there is no existing evidence which can refute the claim that it is s'. Given this, we will need someone claiming s' to present evidence to rule out the possibility that the true state is s. With this in mind, let $B(s)$ denote the set of messages in $M(s)$ which are not elements of any $M(s')$ for which $M(s') \subset M(s)$. Then the burden of proof required to claim s is that evidence must be presented from $B(s)$.

The key assumption we will use regarding the nature of evidence is what Lipman–Seppi call *refutability*. The assumption is that for all s and $s' \neq s$ with $M(s) \not\subset M(s')$, we have $B(s) \subseteq M(s')$. Below, I give some examples to demonstrate that refutability is a weak assumption.

The result: Assume refutability holds. Then if the observer uses the BUR rule, every equilibrium among the debators leads to the observer inferring the true state.

To see why this is true, suppose to the contrary that there is an equilibrium in which in state s, the observer ends up inferring $s' \neq s$. So consider the equilibrium play in state s. By conflicting preferences, there is some debator, say i, who is strictly better off if the observer infers s than when he infers s'. By assumption, i must have a turn to speak. Suppose s'' is the conditionally accepted claim at his turn. Since s'' could not be conditionally accepted unless every state with a nested evidence set has been ruled out, it must be true that $M(s) \not\subset M(s'')$. Hence refutability implies that $B(s) \nsubseteq M(s'')$. Hence i could present evidence in $B(s) \setminus M(s'')$ and claim s. This would refute the conditionally accepted claim of s'' and would satisfy the burden of proof for claiming s. Hence it would make s the conditionally accepted claim. Since s is true, no debator can later refute it. Hence s will be the observer's inference. Since i prefers this to the hypothesized equilibrium outcome, we see that it could not have been an equilibrium. Since a pure strategy equilibrium must exist (since the debators are playing a finite game of perfect information), we see that the observer will end up inferring the correct state.

To see that refutability is a relatively weak assumption, consider the following examples. First, suppose that a state is a realization of N different signals. Suppose the evidence available is that any debator can demonstrate any k signal realizations for some fixed k or range of k's.[6] It is easy to show that refutability holds. Second, suppose that we can order the states as $\{s_1, \ldots, s_n\}$ such that debators can always prove a "lower

[6]This assumption is used by Fishman and Hagerty [1990] and a variant is used by Glazer and Rubinstein [2001].

bound" on the state. That is, if s_j is true, then for any $i \leq j$, there is evidence which proves that the true state is in $\{s_i, s_{i+1}, \ldots, s_n\}$. Formally, $M(s_i) = \{m_1, \ldots, m_i\}$ for each i. For example, it may be that debators can show certain profits earned but cannot prove there are not other profits hidden.[7] Again, it is not hard to show that refutability is satisfied.

In short, BUR rules are surprisingly powerful. Put differently, conditionally accepting statements until they are later refuted provides a strong inducement for truth telling.

Glazer–Rubinstein [2001] (summarized in Chapter 3 of Rubinstein's book) put more structure on the problem and consider situations where complete revelation of information is not possible. When complete revelation is impossible, it is easier to characterize optimal inference rules, *i.e.*, rules which induce the maximum amount of relevation. When full revelation is possible, there are generally too many ways to generate full revelation to enable one to give an interesting characterization.

In their model, a state is a 5–tuple of 1's and 2's. The interpretation is that if the ith component is j, then this supports debator j's position. The observer wishes to learn whether a majority of the aspects favors debator 1 or debator 2. This is all the debators care about as well — a more specific inference is not relevant. The set of feasible messages for debator i in state s is the subset of $\{1, \ldots, 5\}$ consisting of those aspects favoring his position. Given any messages, the observer must conclude in favor of debator 1 or 2. Glazer–Rubinstein characterize inference rules for the observer which minimize the probability he draws a mistaken conclusion.

As Glazer–Rubinstein note, if we add to this model the ability to make a claim of a state, it is easy to get a separating equilibrium. It would simply follow the approach in the example above: the first debator would say which aspects favor him (with or without evidence). If the second debator can refute any of these, the observer concludes in 2's favor. If not, the first debator wins.

Without such messages, one can do part of this. Specifically, suppose we interpret an aspect presented by 1 as a claim of three aspects which are in his favor. For example, providing aspect 1 might be interpreted as a claim that in fact aspects 1, 2, and 3 favor him. If so, 2 must present aspect 2 or 3 to win. If we could do this for all possible three–tuples of aspects, we would achieve separation, just as above. However, there are $\binom{5}{3} = 10$ three–tuples and only 5 aspects. Hence it is impossible to do this perfectly.

There are at least two ways we could think about a partial version of this approach. One way would be to use aspects to encode only 5 of the 10 possible three–tuples. For

[7] This assumption on evidence is used by Okuno–Fujiwara, Postlewaite, and Suzumura [1990] and Nosal [1998].

example, 1 could be interpreted as a claim of 1, 2, and 3, 2 could be interpreted as 2, 3, and 4, etc. This would leave at least 5 states in which the observer reaches the wrong conclusion, namely the 5 three–tuples this misses. An alternative would be to let one aspect encode a claim of more than one three–tuple. For example, aspect 1 could be interpreted as a claim of 1, 2, and 3 *or* 1, 2, and 4 *or* 1, 2, and 5. In this case, refuting 1's claim would require debator 2 to provide aspect 2. Note that the only way to make such a multiple claim work is if there is one aspect (other than the one presented by the first debator) which is common to all the claims. Otherwise, the claim cannot be refuted. Of course, a drawback of this approach is that if only aspects 1 and 2 support debator 1, the observer will mistakenly conclude in debator 1's favor. Glazer–Rubinstein show that the optimal rule combines these two approaches: three aspects are interpreted as claims of three–tuples and the other two as claims of three three–tuples each. If the specific claims are chosen optimally, this leads the observer to err in three out of 32 states.

Glazer–Rubinstein emphasize the asymmetric nature of the optimal inference rule. In particular, it could be that if 1 presents aspect i in his favor, then 2 can win by presenting j as a counterargument, while if 1 presented j, 2 could win by presenting i. As they note, such asymmetries are quite intuitive when one thinks in terms of claims and refutations. If aspect i is used for the initial claim and j for the counterargument, this is informationally very different than when j is used for the claim and i for the counterargument. Put differently, debator 1's message carries much more content than simply the direct implication of the evidence.

There are a number of other papers on communication games with evidence, though none which focus as much on the nature of inference as these two. See, for example, Okuno–Fujiwara, Postlewaite, and Suzumura [1990], Fishman and Hagerty [1990], Shin [1994a, 1994b], Bull and Watson [2001], and Deneckere and Severinov [2001].

Spiegler [2002b] takes a different and quite fascinating approach to modeling inference in debate. He considers multi–issue debates and gives an axiomatic characterization of outcome rules as a function of the available evidence. More specifically, there are two issues, say x and y, each of which could be resolved in either of two ways, giving four possible combinations which I will denote (x, y) (yes on both), $(\neg x, y)$ (no on x and yes on y), etc. Each of the four combinations has a set of attributes associated with it. Spiegler considers debates over each issue separately and a multi–issue debate between (x, y) and $(\neg x, \neg y)$.

Spiegler contrasts positive versus negative arguments. By positive arguments, he means that the relevant evidence is what the various options do satisfy. That is, the set of positive arguments for (x, y) is the set of attributes this position has. The set of positive arguments for x is the set of attributes of the two positions this is consistent with, (x, y) and $(x, \neg y)$. Similarly, the set of positive arguments for $\neg x$ is the set of attributes for $(\neg x, y)$ together with those for $(\neg x, \neg y)$.

The negative arguments for a position consist of those attributes satisfied by the negation of the opponent's position. In one–issue debates, the distinction between positive and negative arguments is meaningless. The negative arguments for x are the positive ones for the negation of $\neg x$ — that is, the positive arguments for x. However, with two–issue debates, the distinction is important. The negative arguments for (x, y) consists of the set of attributes satisfied by the negation of $(\neg x, \neg y)$ — that is, the set of attributes satisfied by $(\neg x, y)$, $(x, \neg y)$, and (x, y).

A resolution function for a debate takes the relevant set of arguments (either positive or negative, depending on which case we are considering) and decides a winner. Note that it is purely a weighing of arguments — it cannot depend on whether the debate is single issue or multi–issue. Spiegler considers two simple axioms on this function. First, he assumes that if, given a particular assignment of attributes, x would beat $\neg x$ and y would beat $\neg y$, then it must be true that (x, y) beats $(\neg x, \neg y)$. This consistency across debate formats, termed *procedural invariance* by Spiegler, seems quite reasonable. Second, he assumes that if, say, debator 1 wins with a particular attribute assignment, then he still wins if we enlarge the set of attributes in his favor and/or shrink the set of attributes in the other debator's favor. If we think of the debators as choosing which arguments to present and view the resolution function as a prediction regarding the equilibrium outcome, then this seems hard to argue with. Spiegler calls this axiom *free disposal*.

Spiegler's results are intriguing. First, he shows that with positive arguments, the only resolution functions consistent with his axioms are constant ones. That is, debate is trivial. To see the intuition, suppose the set of all attributes is $\{A, B\}$. Let r be the resolution function and suppose it is not constant. So $r(C, D) = 1$ means that 1 wins when the arguments on his side are those in the set C and the arguments for 2 are those in the set D. I will abbreviate $\{A, B\}$ as AB below.

If $r(A, AB) = r(B, AB) = 1$, then free disposal implies that r is constant at 1. Hence, without loss of generality, suppose $r(A, AB) = 2$. If $r(AB, A) = r(AB, B) = 2$, then 2 always wins. So either $r(AB, A) = 1$ or $r(AB, B) = 1$. Suppose it is the former. Since $r(A, AB) = 2$, free disposal implies $r(A, A) = 2$. But $r(AB, A) = 1$ and free disposal implies the opposite. Hence we must have $r(AB, B) = 1$ and so $r(A, B) = 1$ by free disposal. Summarizing, without loss of generality, we can assume that $r(A, AB) = 2$ and $r(A, B) = 1$.

Consider the attribute assignment where the only attribute for each of (x, y), $(\neg x, y)$ and $(x, \neg y)$ is A, while the only attribute of $(\neg x, \neg y)$ is B. In this case, the set of positive arguments for x is just A, while the positive arguments for $\neg x$ is $\{A, B\}$. Similarly, the only positive argument for y is A, while the positive arguments for $\neg y$ are $\{A, B\}$. Given the assumed resolution function, then, 2 wins both of these debates. By procedural invariance, 2 must win the multi–issue debate as well. But the only positive argument for (x, y) is A, while the only positive argument for $(\neg x, \neg y)$ is B. Given the assumed resolution function, 1 wins.

Interestingly, Spiegler shows that this problem does not arise for negative arguments. To see the point, return to the attribute assignment and resolution function discussed above. The negative arguments in the one issue debate are the same as the positive ones, so, as before, 2 would win both one issue debates. The only negative argument for (x, y) is A, while the negative arguments for $(\neg x, \neg y)$ are $\{A, B\}$. Hence 2 wins the multi–issue debate as well, just as procedural invariance requires.

An intriguing property of the resolution function in the case of negative arguments is a kind of bias in favor of one side or the other. Specifically, these rules have the property that there is some set of attributes such that one particular agent (say 1) wins whenever his set of arguments contains *any* of these attributes. In this sense, 1 has a favored position. Spiegler interprets this as a status quo bias.

Spiegler suggests that his approach may help make sense of rhetorical arguments in real debates. For example, his results indicate that in a debate between (x, y) and $(\neg x, \neg y)$, changes in the attributes of $(x, \neg y)$ — a position neither debator advocates — can change the outcome of the debate. Arguably, we do see such apparently irrational reactions.

5 Vagueness

The last topic I wish to discuss[8] is one which, to my knowledge, has been ignored entirely by the economics literature: vagueness[9]. When one thinks about language as spoken by people on a day–to–day basis, it is hard to ignore the fact that much of what is said is vague. Consider, for example, the word "tall." There is no precise, known height which defines the line between a person who is tall and a person who is not. Why do we use a language in which such terms are so prevalent? Why don't we simply adopt as a definition that "tall" will mean above, say, 6 foot 2? We could even adopt a context–specific definition, saying for example that "tall" for a newborn means above 15 inches, while "tall" for a professional basketball player means above 6 foot 10.

In this section, I will argue that we cannot explain the prevalence of vague terms in natural language without a model of bounded rationality which is significantly different

[8]This section is based on Lipman [2001].

[9]Keefe and Smith [1996] is an excellent introduction to the philosophy literature on this topic.

from anything (that I know) in the existing literature. In a nutshell, the argument is that any model along existing lines will imply that a precise language like the one described above would Pareto dominate the vague language we see in every society in history. Of course, it seems rather far–fetched to conclude that we have simply tolerated a world–wide, several–thousand–year efficiency loss. Further, even a moment's reflection will suggest that it is easier to speak when one is allowed to use vague language than it would be if such language were banned. Hence this dominance surely tells us that there is something wrong with the model, not the world.

First, let me try to be more precise about what I mean by vagueness. Following Sainsbury [1990], I will say that a word is *precise* if it describes a well–defined set of objects. By contrast, a word is *vague* if it is not precise. Hence the alternative definition given above for "tall" would make this term precise, whereas it is vague in its current usage.

I emphasize that vagueness (as I use the term) is *not* the same thing as less than full information. To say that a person's height is above six feet is precise in the sense that it defines a set of people unambiguously. It is less informative than saying that a person's height is 6 foot 2, but it is *not* vague as I use the term.

The classic illustration of vagueness is what is referred to in the philosophy literature as the *sorites paradox*. For one version of this paradox, note that two facts seem clear about the way people use the word "tall." First, anyone whose height is 10 feet is tall. Second, if one person is tall and a second person's height is within 1/1000 of an inch of the first, then the second person is tall as well. But then working backward from 10 feet, we eventually reach the absurd conclusion that a person whose height is 1 inch is tall. Of course, the source of the difficulty is the vagueness of the word "tall." Because there is no fixed set of heights which "tall" corresponds to, no line which separates "tall" from not, one has the feeling that a small change should not matter. Of course, many small changes add up to a big change, so this is not consistent.

Many other words are vague in this sense. Some words which can easily yield a sorites paradox are "bald," "red" (imagine a sequence of objects moving continuously from red to orange to yellow), "thin," "child" (imagine a sequence of people, each one second older than the previous), "many," and "probably." Some philosophers have constructed less obvious sorites paradoxes for "tadpole," "chair," and other seemingly clear–cut terms.

The prevalence of vague terms in natural language poses two intriguing and interrelated challenges. First, what meaning do such terms convey? Second, why are they so prevalent?

As to the first question, it seems clear that vague terms acquire their meaning from usage. That is, whatever meaning "tall" has is due to the way people use the word, not any particular logical structure. The most obvious way to model this formally is to treat vague terms as ones which are used probabilistically, so the probability someone is described as "tall" is increasing in height but is strictly between 0 and 1 for a certain range. This kind of uncertainty could correspond to mixed strategies or private information, either of which would give a clear notion of the meaning of a vague term.

However, this approach has a severe drawback: it cannot give a good answer to the second question. In particular, if this is what vagueness is, we would be better off with a language which replaced vague terms with precise ones. To be sure, many vague terms would be difficult to redefine in a precise way. On the other hand, many can easily be given such redefinitions. As noted above, "tall" could be defined to be over a specific height, "bald" could be defined by a specific fraction of the skull covered by hair, etc.[10] In such cases, if vagueness is to be interpreted as a probability distribution over interpretations, we would be better off using precise definitions.

To see the point, consider the following simple example. Player 2 must pick up Mr. X at the airport but has never met him. Player 1, who knows Mr. X, can describe him to 2. Suppose that the only variable which distinguishes people is height and that this is independently distributed across people uniformly on $[0, 1]$. 1 knows the exact height of Mr. X; 2 does not. However, both know that when 2 gets to the airport, there will be three people there, Mr. X and two (randomly chosen) others. 2 has very little time, so he can only ask one person if he is Mr. X. If 2 chooses correctly, 1 and 2 both get a payoff of 1; otherwise, they both get 0. Clearly, there is a simple solution if 2 can observe the heights of the people at the airport and the set of possible descriptions is $[0, 1]$: 1 can tell 2 Mr. X's height. Because the probability that two people are the same height is zero, this guarantees that Mr. X will be picked up.

However, this is a much bigger language than any in use. So let us take the opposite extreme: the only descriptions 1 can give are "short" and "tall." Further, 2 cannot observe the exact height of any of the people at the airport, only relative heights. That is, 2 can tell who is tallest, who is shortest, and who is in the middle.

In this case, it is not hard to show what the efficient language is: 1 should say "tall" if Mr. X's height is greater than $1/2$ and "short" otherwise.[11] 2 tries the tallest person at the airport when he is told that Mr. X is "tall" and the shortest when told that Mr. X is "short." Note, in particular, that there is no vagueness in the optimal language: "tall" corresponds to the set $[1/2, 1]$.

[10] Of course, sometimes we do develop such definitions for legal purposes, most notably in the case of the word "child." On the other hand, it is clear that common usage of this word outside the legal context is not based on such a precise definition.

[11] Of course, the words themselves are not relevant to this equilibrium. An equally efficient language would reverse the roles of "tall" and "short" or even replace them with "middle" and "blond."

What would vagueness mean in this context? One way to make "tall" vague would be if 1 randomizes. For example, suppose 1 says "short" if Mr. X's height is below 1/3, "tall" if it is greater than 2/3, and randomizes in between with the probability he says "tall" increasing with Mr. X's height. While this resembles the way "tall" is used in reality, there are no equilibria of this kind. If one takes a nonequilibrium approach and assumes 1 is committed to such a language, it is easy to show that both 1 and 2 (and presumably Mr. X!) would be better off in the pure strategy equilibrium above.

Alternatively, private information could give the needed randomness. Suppose 1 has observed some signal in addition to height. In this case, the efficient language partitions not the set of heights but the set of height–signal pairs. Hence a word will correspond to a random statement about height, the randomness being induced by the signal. On the other hand, what is this other signal? First, suppose it is somehow intrinsically relevant. For example, 1 may know Mr. X's weight and 2 may be able to observe relative weights. In this case, 1 needs to communicate on two dimensions to 2. The efficient language will have terms which are precise in two dimensions even though this may make them imprecise in any one dimension. This does not seem to be much of an explanation of an apparently unidimensional term like "tall." On the other hand, suppose the signal is not relevant. Then it would be most efficient to ignore this signal and use the language described above.

Why, then, are vague terms so prevalent? There are several seemingly obvious answers to this question which closer examination calls into doubt. For example, one obvious advantage to vague language is that it makes context–sensitivity easier. If "tall" is not given a precise definition, I can use it to describe a newborn whose height is 2 feet or a professional basketball player whose height is 7 feet. This answer doesn't work, however: we could make the precise definitions context–specific as suggested above. That is, "tall" could be defined to be greater than or equal to 15 inches for a newborn and greater than or equal to 6 foot 10 for a professional basketball player. In terms of the model above, it would be simple to add a random variable which is observed by both players and interpreted as the context. The result above would imply that in each context, it is optimal to have a precise language, though the language might vary with the context.

A natural objection to this point is that it is cognitively difficult to remember all the relevant cutoff points. Of course, the key question is not whether it is difficult to remember the cutoff points corresponding to words in a precise language but whether it is more difficult than remembering a vague meaning. Consideration of such a question requires a model of bounded rationality different from anything in the existing literature.

One explanation for the prevalence of vague terms which has been suggested in the philosophy literature is that vague terms might be easier to learn. Sainsbury [1990] argues that "we acquire the concept from the inside, working outwards from central cases ... rather than from the outside, identifying boundaries and moving inwards." As I understand the argument, the idea is rather natural. Suppose that I enter a population of people who have agreed on a precise cutoff height between "short" and "tall." The only way I can learn about this cutoff is to observe how they classify various people. Naturally, unless I get very lucky and observe the classification of two individuals very close to but on opposite sides of the boundary, I will never learn the cutoff being used. Instead, I will end up learning that people below, say, 5 foot 5, are short, while people above, say, 6 foot 6, are tall and I would have no clear idea about people in the middle. In this sense, what I have learned is a vague term. That is, even if we begin with a precise term, as new agents come into the population and old ones die out, the term will become vague.

While this idea has some appeal for terms based on variables that are difficult to quantify (such as "nice" or perhaps Sainsbury's example of "red"), it is hard to accept when applied to a notion like height. Surely in the fictitious world described above, I should be able to open a dictionary and get out a tape measure to see exactly how tall is "tall."[12]

I think the only way to formally understand the prevalence of vague terms is in a model with a different kind of bounded rationality than what is considered in the literature. There are at least three possibilities which I give in increasing order of ambitiousness. First, vagueness may be easier than precision, for the speaker, listener, or both. For the speaker, deciding which precise term to use may be harder than being vague. For the listener, information which is too specific may require more effort to analyze. With vague language, perhaps one can communicate the "big picture" more easily. This requires a different model of information processing than any I know of.

A more difficult approach would be to derive vagueness from unforeseen contingencies. If the speaker does not know all the possible situations where the listener would use the conveyed information, it may be optimal to be vague. For example, contracts often use vague terms such as "taking appropriate care" or "with all due speed" instead of specifying precisely what each party should do. If agents fear that circumstances may arise that they have not yet imagined, then they may avoid precision to retain flexibility. Hence the optimal contract may require the parties to respond to unexpected circumstances "appropriately," with the hope that the meaning of this word will be sufficiently clear *ex post*.[13] Perhaps a similar phenomenon also leads to vagueness in language outside the

[12] Another problem is that it is hard to see why the language won't ultimately collapse. After all, each generation learns a vague version of the previous generation's usage. For example, one can give processes for which the language converges to a distribution of cutoffs which is uniform on the set of heights, as uninformative a language as one could imagine.

[13] This idea is very similar to the Grossman–Hart–Moore approach to incomplete contracts. See Hart [1995] on this approach and Dekel, Lipman, and Rustichini [1998] for a discussion of the connection between it and formal models of unforeseen contingencies.

contractual setting. Given the difficulty of modeling unforeseen contingencies (see Dekel, Lipman, and Rustichini [1998] for a survey), this approach is surely not easy.

Finally, I turn to a still more ambitious approach. To motivate it, consider one seemingly obvious reason why vague terms are useful: the speaker might not observe the height of an individual well enough to be sure how to classify him precisely. If we modify the example to include this, however, 1 would have subjective beliefs about Mr. X's height and the efficient language would partition the set of such probability distributions. Hence this objection simply shifts the issue: why don't we have a precise language for describing such distributions?

An obvious reply is that real people do not form precise subjective beliefs. Again, though, this is not a sufficient objection. Take your favorite model of imprecise beliefs — belief functions, sets of priors, nonadditive probabilities, whatever. Then surely the optimal language corresponds to a precise partition of the set of imprecise beliefs. Hence the issue of whether beliefs are precise or not is irrelevant; the real question is whether people have a belief in mind at all.

If one takes the Savage view of subjective beliefs, one must interpret this as saying that agents do not have preferences or, perhaps, that agents do not truly "know" their own preferences. If we think of preferences over, say, flavors of ice cream, this sounds ridiculous. If we think of preferences over state–contingent sequences of commodity bundles over one's lifetime, it seems obviously correct. In 1967, Savage described the problem as "the utter impracticality of knowing our own minds in the sense implied by the theory." He went on to comment

> You cannot be confident of having composed the ten word telegram that suits you best, though the list of possibilities is finite and vast numbers of possibilities can be eliminated immediately; after your best efforts, someone may suggest a clear improvement that had not occurred to you.

Put differently, the vastness of even very simple sets of options suggests it is ludicrous to think a real person would have a meaningful understanding of the space, much less well behaved preferences over it.

In short, it is not that people have a precise view of the world but communicate it vaguely; instead, they have a vague view of the world. I know of no model which formalizes this.

References

[1] Blume, A., Y.-G. Kim, and J. Sobel, "Evolutionary Stability in Games of Communication," *Games and Economic Behavior*, 5, 1993, 547–575.

[2] Blume, A., "A Learning–Efficiency Explanation of Structure in Communication," working paper, 1999.

[3] Blume, A., "Coordination and Learning with a Partial Language," *Journal of Economic Theory*, 95, November 2000, 1–36.

[4] Bull, J., and J. Watson, "Evidence Disclosure and Verifiability," working paper, 2001.

[5] Crawford, V., and J. Sobel, "Strategic Information Transmission," *Econometrica*, 50, November 1982, 1431–1451.

[6] Dekel, E., B. Lipman, and A. Rustichini, "Recent Developments in Modeling Unforeseen Contingencies," *European Economic Review*, 42, 1998, 523–542.

[7] Deneckere, R., and S. Severinov, "Mechanism Design and Communication Costs," working paper, 2001.

[8] Farrell, J., "Meaning and Credibility in Cheap Talk Games," *Games and Economic Behavior*, 5, 1993, 514–531.

[9] Fishman, M., and K. Hagerty, "The Optimal Amount of Discretion to Allow in Disclosure, *Quarterly Journal of Economics*, 1990, 427–444.

[10] Glazer, J., and A. Rubinstein, "Debates and Decisions: On a Rationale of Argumentation Rules," *Games and Economic Behavior*, 36, August 2001, 158–173.

[11] Hart, O., *Firms, Contracts, and Financial Structure*, Oxford: Clarendon Press, 1995.

[12] Keefe, R., and P. Smith, *Vagueness: A Reader*, Cambridge, MA: MIT Press, 1996.

[13] Lipman, B., and D. Seppi, "Robust Inference in Communication Games with Partial Provability," *Journal of Economic Theory*, 66, August 1995, 370-405.

[14] Lipman, B., "Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality," *Review of Economic Studies*, 66, April 1999, 339-361.

[15] Lipman, B., "Why is Language Vague?," working paper, 2001.

[16] Nosal, E., "Financial Distress and Underemployment," *Review of Economic Studies*, 65, October 1998, 817–845.

[17] Okuno–Fujiwara, M., A. Postlewaite, and K. Suzumura, "Strategic Information Revelation," *Review of Economic Studies*, 57, 1990, 25–47.

[18] Rubinstein, A., *Economics and Language*, Cambridge: Cambridge University Press, 2000.

[19] Sainsbury, R. M., "Concepts without Boundaries," Inaugural Lecture, King's College, 1990. Reprinted in *Vagueness: A Reader*, ed. by R. Keefe, and P. Smith, Cambridge, MA: MIT Press, 1996, 251–264.

[20] Savage, L. J., "Difficulties in the Theory of Personal Probability," *Philosophy of Science*, 34, 1967, 305–310.

[21] Shafir, E. , I. Simonson, and A. Tversky, "Reason–Based Choice," *Cognition*, 49, 1993. Reprinted in *Choices, Values, and Frames*, ed. by D. Kahneman and A. Tversky, Cambridge: Cambridge University Press, 2000, 597–619.

[22] Shin, H. S., "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, 64, 1994a, 253–263.

[23] Shin, H. S., "News Management and the Value of Firms," *Rand Journal of Economics*, 25, 1994b, 58–71.

[24] Spence, A. M., *Market Signaling*, Cambridge, MA: Harvard University Press, 1974.

[25] Spiegler, R., "Equilibrium in Justifiable Strategies," *Review of Economic Studies*, forthcoming.

[26] Spiegler, R., "Justifiability and Delay Tactics in a Concession Game," Tel Aviv University working paper, 2002a.

[27] Spiegler, R., "Argumentation in Multi–Issue Debates," working paper, 2002b.

[28] Tversky, A., and D. Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science*, 211, 1981, 453–458.

[29] Wärneryd, K., "Cheap Talk, Coordination, and Evolutionary Stability," *Games and Economic Behavior*, 5, 1993, 532–546.