

Comments on Rubinstein's *Economics and
Language*¹

Barton L. Lipman

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

email: blipman@ssc.wisc.edu

September 1999

¹I thank Marie–Odile Yanelle for helpful conversations. No thanks are due to the Dell Corporation.

“Is he — is he a tall man?”

“Who shall answer that question?” cried Emma. “My father would say, ‘Yes’; Mr. Knightly, ‘No’; and Miss Bates and I, that he is just the happy medium.”

— From *Emma* by Jane Austen.

1 Introduction

While a reader of this book may be surprised to see a game theorist writing about language, he should instead be surprised by how few game theorists have done so. As Rubinstein observes, language is a game: I make a statement because I believe you will interpret it in a particular way. You interpret my statement based on your beliefs about my intentions. Hence speaker and listener are engaged in a game which determines the meaning of the statement.

Furthermore, I think more “traditional” economists should be interested in models of language. 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.

Rubinstein gives an intriguing opening to this important topic. In what follows, I give a brief overview of what I see as the important ideas in the book and use this as background to comments on directions for future research. In the process, I will suggest that there are important and difficult problems ahead.

2 Overview

While Rubinstein jokes that there is little to connect his lectures beyond the common title and speaker, I think it’s conceptually useful to consider the various ingredients of a theory of language the book gives. As I see it, the first four chapters give us various ways to address three questions:¹

¹I have only one comment on Chapter 5. Rubinstein criticizes the way game theory is “sold” to business students, saying that the kind of conceptual understanding that game theory can provide

- What gives a statement its meaning?
- Why is language the way it is?
- How does language affect actions?

Rubinstein gives two distinct ways to answer the first question. Chapters 1 and 4 give one of his approaches, modeling language as a *logical system*. In this approach, language is a set of building blocks and rules of construction which implicitly define the meaning of statements. For example, the meaning of $p \vee q$ is that either p or q is true — a meaning which, of course, depends on the underlying meaning of p and q .

While this is a natural starting point, it is clear that language as spoken by real people is *not* logical in the formal sense. For example, as Rubinstein notes in Chapter 3, a listener generally makes inferences in response to a statement which go beyond the purely logical implications of what the speaker said. Put differently, the fact that a given statement is made can itself signal some of the speaker’s private information. Later, I will discuss another way in which language is not purely logical.

Rubinstein’s second way to answer question 1 addresses this drawback of the logical system approach, adopting instead what I’ll call the *equilibrium approach*. This approach, used in Chapters 2 and 3, treats language as a set of words and an equilibrium interpretation of their usage or meaning. For example, if a particular statement is only made in certain situations, then the meaning of the statement (in addition to any concrete evidence included, as studied in Chapter 3) is that one of these situations must be true. While the equilibrium approach captures some of what the logical system approach misses, it does so at a cost: the language has no structure to it. Hence many questions about language cannot be addressed in such a model.

Rubinstein’s approach to the second question, illustrated nicely in Chapters 1 and 3, is what I’ll call the *structural optimization hypothesis*. In this approach, language is that structure which maximizes the amount of information conveyed subject to constraints on the “complexity” of the language. In Chapter 1, the “complexity constraint” is that

“stop[s] far short of giving advice and claiming to be useful.” I agree up to a point. Having taught game theory to MBA students for a few years, I spent much time reflecting on what value, if any, the theory had for them. I think most business students want rules, not concepts, evidently believing that the business world is a series of clearly framed problems, each with an appropriate rule to follow. This absurd notion influences the way textbooks are written since these are supposed to appeal to business students. I disagree with Rubinstein, though, when he concludes that there is nothing “useful” to such people in game theory. I think it is precisely the concepts and understanding he alludes to which can benefit someone in business. Of course, their understanding will never approach Rubinstein’s! In particular, backward induction is indeed a revelation to business students and the concept is hard for them to learn. Any reader who is skeptical on this point need only see how many business students will err on a question involving sunk costs.

the language can only have one relation. The question then is what relations are most useful for describing many objects (or sets of objects) as precisely as possible. The complexity constraint in Chapter 3 is that the listener cannot process more than two pieces of evidence. Here Rubinstein considers the optimal way for the listener to interpret statements by debaters in order to elicit information from them.

Finally, Chapter 4 gives an intriguing approach to answering the third question which I'll refer to as the *expressibility effect*. The general perspective here is that people perceive the world and form decisions in words. Because of this, the nature of the language people use affects their actions. Chapter 4 looks at the expressibility effect through the requirement that decision rules or preferences be definable, but, as Rubinstein notes, the principle could be explored in many other ways and other aspects of decision making.

3 Further Research

I now turn to three areas for future research in this area which strike me as particularly interesting. First, I am intrigued by the inherent circularity one gets when combining the structural optimization hypothesis and the expressibility effect. By means of illustration, one response to Chapter 1 which I have heard is that it may be that linear orderings are common in our language not because of any inherent usefulness they have but simply because linear orderings are common in the natural world. I think this criticism misses a crucial point: do we perceive linear orderings to be common in the world because they are common in our language? In other words, the structural optimization approach suggests that we structure language in a way which seems useful to us given our perception of the world. On the other hand, because of the expressibility effect, once developed, language affects the way we see the world. An analysis which includes such feedback effects could be quite interesting.

Second, as noted earlier, both the logical systems and equilibrium approaches to meaning miss part of the picture. It is clear that language does have structure even if it is not fully logical. It is not obvious how to model this simple notion. If the meaning of a sentence comes from its equilibrium interpretation only, then there is no reason for structure to relate to content. "I live in Wisconsin and it is cold" could be interpreted as the conjunction of "I live in Wisconsin" and "It is cold" or the disjunction of "I am hungry" and "Rubinstein lives in Israel." At the same time, meaning cannot come purely from the content of the sentence in isolation from its context or other extralogical factors. A model which combines the advantages of equilibrium models and logical systems would be more plausible and, perhaps, more useful.

It might be possible to develop such a model using the structural optimization ap-

proach. In this case, a particular natural complexity constraint is limited memory. It seems intuitively obvious that it is easier to remember that “I live in Wisconsin and it is cold” means “I live in Wisconsin” and “It is cold” instead of some notion unrelated to Wisconsin or cold. Similarly, the structure of words might be derivable from memory limitations. For example, consider the use of prefixes and suffixes. To deduce the meaning of a word beginning with “post,” it is sufficient to remember the meaning of the prefix “post” and the word which follows. To combine this with an equilibrium approach to meaning, one would presumably focus on “efficient” equilibria.

The last topic is another aspect of the issue of meaning: vagueness.² Perhaps it is easier to understand what I mean by “vague” by contrasting it with precise. I will say a term is *precise* if it describes a well-defined set of objects. This is the way language works in most of the economic models where it appears: the set of objects is partitioned and a word is associated with each event of the partition. By contrast, a term is *vague* if it does not identify such a set.

To illustrate my meaning and to show why vague terms make it difficult to model language as a logical system, consider the following version of the famous *sorites paradox*. Two facts seem clear regarding the way most 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 clear: “tall” does not correspond to a clearly defined set. There is no fixed height which defines the line between someone who is tall and someone who is not. Because of this, there is an inherent ambiguity in the meaning of the word “tall” as the quote I began with illustrates. Many other words have this property: consider “bald,” “red,” “thin,” “child,” “many,” and “probably.”

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. Hence it seems natural to follow the equilibrium approach to studying meaning. The most obvious way to do so 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.

²Keefe and Smith [1996] is a fascinating introduction to the philosophy literature on the subject.

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 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.³ 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 $[1/2, 1]$.

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.⁴

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.

³Of 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.”

⁴Lipman [1999] gives a generalization of this argument and some further discussion.

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? It seems rather obvious that such words are useful — a moment’s reflection will suggest that it would be difficult to say much if one were not allowed to be vague!

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 this book or in the literature.⁵ 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*.⁶ 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?

⁵I give a fuller defense of this position in Lipman [1999].

⁶This 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.

An obvious reply is that real people do not form precise subjective beliefs. In other words, 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.⁷ I know of no model which formalizes this. For example, it is not enough to replace probability distributions with nonadditive probabilities. If agents have nonadditive probabilities, then it is surely optimal to partition the set of such beliefs precisely.

To sum up, I think this book gives some intriguing starts on some important problems. As such, it opens the door to some exciting and difficult research.

References

- [1] Dekel, E., B. Lipman, and A. Rustichini, “Recent Developments in Modeling Unforeseen Contingencies,” *European Economic Review*, 42, 1998, 523–542.
- [2] Hart, O., *Firms, Contracts, and Financial Structure*, Oxford: Clarendon Press, 1995.
- [3] Keefe, R., and P. Smith, *Vagueness: A Reader*, Cambridge, MA: MIT Press, 1996.
- [4] Lipman, B., “Why is Language Vague?,” working paper, 1999.

⁷The dividing line between unforeseen contingencies and this kind of vague perception is itself quite vague!