

THE CHURCHILL LECTURES IN ECONOMIC THEORY

Economics and Language

Five Essays

ARIEL RUBINSTEIN

Tel Aviv University

Princeton University



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK www.cup.cam.ac.uk
40 West 20th Street, New York, NY 10011-4211, USA www.cup.org
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain

© Ariel Rubinstein 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000

Printed in the United Kingdom at the University Press, Cambridge

Typeface Trump Mediaeval 9.75/12 pt System QuarkXPress® [SE]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Rubinstein, Ariel.
Economics and language / Ariel Rubinstein.
p. cm.
Includes bibliographical references.
ISBN 0-521-59306-9 (hardbound)
1. Economics—Language. 2. Game theory. I. Title.
HB62.R83 2000
330'.01'4—dc21
ISBN 0 521 593069 hardback
ISBN 0 521 789907 paperback

CONTENTS

<i>Acknowledgments</i>	<i>page viii</i>
------------------------	------------------

PART 1 ECONOMICS OF LANGUAGE

0 Economics and language	3
1 Choosing the semantic properties of language	9
2 Evolution gives meaning to language	25
3 Strategic considerations in pragmatics	37

PART 2 LANGUAGE OF ECONOMICS

4 Decision making and language	55
5 On the rhetoric of game theory	71
Concluding remarks	89

PART 3 COMMENTS

Johan van Benthem	93
Tilman Börgers	108
Barton Lipman	114
<i>Index</i>	125

CHAPTER 1

CHOOSING THE SEMANTIC PROPERTIES OF LANGUAGE

1.1 Introduction

This chapter will present a research agenda whose prime objective is to explain how features of natural language are consistent with the optimization of certain “reasonable” target functions. Rather than discuss the research agenda in abstract, I will begin with the specific argument and return to the general discussion at the end of the chapter.

This chapter discusses binary relations. A binary relation on a set Ω specifies a connection between elements within the set. Such binary relations are common in natural language. For example, “person x knows person y ,” “tree x is to the right of tree y ,” “picture x is similar to picture y ,” “chair x and chair y are the same color,” and so on. I will avoid binary relations such as “Professor x works for university y ” or “the Social Security number of x is y ,” which specify “relationships” between elements which naturally belong to two distinct sets. I will further restrict the term “binary relation” to be irreflexive: No element relates to itself. The reason for this is that the term “ x relates to y ” when $x=y$ is fundamentally different from “ x relates to y ” when $x \neq y$. For example, the statement “ a loves b ” is different from the statement “ a loves himself.”

Certain binary relations, by their nature, must satisfy certain properties. For example, the relation “ x is a neighbor of y ” must, in any acceptable use of this relation, satisfy the symmetry property (if x is a neighbor of y , then y is a neighbor of x). The relation “ x is to the right of y ”

This chapter is based on Rubinstein (1996).

must be a *linear ordering*, thus satisfying the properties of completeness (for every $x \neq y$, either x relates to y or y to x), asymmetry (for every x and y , if x relates to y , y does not relate to x), and transitivity (for every x , y , and z , if x relates to y and y to z , then x relates to z). In contrast, the nature of many other binary relations, such as the relation " x loves y ," does not imply any specific properties that the relation must satisfy *a priori*. It may be true that among a particular group of people, " x loves y " implies " y loves x ." However, there is nothing in our understanding of the relation " x loves y " which necessitates this symmetry.

The subject of this chapter is in fact the *properties* of those binary relations which appear in natural language. Formally, a property of the relation R is defined to be a sentence in the language of the calculus of predicates which uses a name for the binary relation R , variable names, connectives, and qualifiers, but does not include any individual names from the set of objects Ω . I will refer to the combination of properties of a term as its *structure*.

I am curious as to the structures of binary relations in natural language. I search for explanations as to why, out of an infinite number of potential properties, we find that only a few are common in natural languages. For example, it is difficult to find natural properties of binary relations such as the following:

- A1: If xRy and xRz ($y \neq z$), and both yRa and zRa , then also xRa .
- A2: For every x there are *three* elements y for which xRy .
(In contrast, the relation " x is the child of y " on the set of human beings does satisfy the property that for every x there are *two* elements y which x relates to.)

Alternatively, it is difficult to find examples of natural structures of binary relations which are required to be tournaments (satisfying completeness and asymmetry) but which are not required to satisfy transitivity. One exception is the structure of the relation " x is located clockwise from y (on the shortest arc connecting x and y)."
Is it simply a coincidence that only a few structures exist in natural language?

Choosing the semantic properties of language

The starting point for the following discussion is that binary relations fulfill certain functions in everyday life. There are many possible criteria for examining the functionality of binary relations. In this discussion, I examine only three. I will argue that certain properties, shared by linear orderings, perform better according to each of these criteria. Of course, other criteria are also likely to provide alternative explanations for the frequent use of various common structures such as equivalence and similarity relations.

1.2 Indication-friendliness

Consider the case in which two parties observe a group of trees and the speaker wishes to refer to a certain tree. If the tree is the only olive tree in the grove, the speaker should simply use the term "the olive tree." If there is no mutually recognized name for the tree and the two parties have a certain binary relation defined on the set of trees in their mutual vocabulary, the user can use this relation to define the element. For example, the phrase "the third tree on the right" indicates one tree out of many by using the linear ordering "x stands to the left of y" when the group of trees is well defined and the relation "being to the left of" is a linear ordering. Similarly, the phrase "the seventh floor" indicates a location in a building given the linear ordering "floor x is above floor y." There would be no need to use the phrase if it was known to be "the presidential floor." On the other hand, the relation "line x on the clock is clockwise to line y (with the smallest angle possible)" does not enable the user to indicate a certain line on a number-less clock; any formula which is satisfied by three o'clock is satisfied by four o'clock as well. In fact, the existence of even one designated line such as "twelve o'clock", would enable the use of the relation to specify all lines on the clock. The effect of using such a designated element is equivalent to transforming the circle into a line.

Thus, binary relations are viewed here as tools for indicating elements in a set whose objects do not have names.

We look for structures that enable the user to unambiguously single out *any element out of any subset of Ω* . We are led to the following definition:

Definition: A binary relation R on a set Ω is *indication-friendly* if for every $A \subseteq \Omega$, and every element $a \in A$, there is a formula $f_{a,A}(x)$ (in the language of the calculus of predicates with one binary relation and without individual constants) such that a is the only element in A satisfying the formula (when substituting a in place of the free variable x).

All linear orderings are indication-friendly. If R is a linear ordering, the formula $P_1(x) = \forall y(x \neq y \rightarrow xRy)$ defines the “maximal” element in the set A for $A \subseteq \Omega$. The formula $P_2(x) = \forall y(x \neq y \wedge \neg P_1(y) \rightarrow xRy)$ defines the “second-to-the-maximal” element, and so on. Note that in natural language there are “short cuts” for describing the various elements. For example, the “short cuts” for $P_1(x)$ and $P_2(x)$ are “the first” and “the second.”

In contrast, consider the set $\Omega = \{a, b, c, d\}$ and the non-linear binary relation R , called “beat,” depicted in the following diagram (aRb , aRc , dRa , bRd , bRc and cRd):




Figure 1.1

Referring to the grand set Ω , the element a is defined by “it beats two elements, one of which also beats two elements.” The element b is defined by “it beats two elements, which each beat one element.” And so on. However, whereas the relation R allows the user to define any element in the set Ω , the relation is not effective in

Choosing the semantic properties of language

defining elements in the subset $\{a,b,d\}$, in which case the induced relation is cyclical.

We will now demonstrate that if Ω is a finite set and R is a binary relation, then R is indication-friendly if and only if R is a linear ordering. We have already noted that if R is a linear ordering on Ω , then for every $A \subseteq \Omega$ and every $a \in A$, there is a formula which indicates a . Assume that a binary relation R is indication-friendly. For any two elements $a,b \in \Omega$, in order to indicate any of the elements in the two-member set $A = \{a,b\}$, it must be that either aRb or bRa but not both; thus, R must be complete and asymmetric. R must *also* be transitive since for every three elements $a,b,c \in \Omega$, the relation must not be cyclical in order to indicate each of the elements in the set $A = \{a,b,c\}$.

Conclusion 1: A binary relation enables the user to indicate any element in any subset of the grand set if and only if it is a linear ordering. Linear orderings are the most efficient binary relations for indicating every element in every subset.

1.3 A detour: splitting a set

We will now make a brief detour from the world of binary relations in order to discuss unary relations. Assume that the user of the language can refer to a set of objects X (such as “the set of flowers”). From time to time, the speaker will wish to refer to subsets of X (either for the purpose of conversing with another person or for storing information in his own mind). However, he will be able to refer only to terms that appear in his language. Initially, the speaker can refer to the set of all X s (“pick all flowers”) or to the null set (“don’t pick any flowers”). In order to extend his vocabulary the designer of the language is permitted to invent (given “hardware” constraints) one additional term for one subset of X . The objective of the designer is to introduce one new term so that the speaker can refer to a new set in as accurate a manner as possible (on average).

Let us be more precise: Restrict X to be a finite set. If the term “the set S ” is well defined, then the user will have

four expressions available for referring to subsets of X : “all X s,” “the S s,” “the not S s,” and “nothing.” We will call the collection of sets $V(S) = \{X, S, -S, \emptyset\}$ the *vocabulary* spanned by S .

The basic idea is that language should be flexible enough to function under unforeseen circumstances. The speaker can use the terms in $V(S)$ but will eventually need to refer to sets not necessarily contained in $V(S)$. The term S will be evaluated by the vocabulary’s ease of use with “least loss.” It is assumed that when the speaker wishes to refer to a set $Z \subseteq X$, he will use an element in $V(S)$ which is “closest” to the set Z . In order to formally state this idea, we need to define the *distance* between two sets. Let the distance between A and B , $d(A, B)$ be the cardinality of the asymmetric difference between A and B (the set of all elements which are in A and not in B or in B and not in A) – i.e. $d(A, B) = |(A - B) \cup (B - A)|$. One interpretation of this distance function fits the case in which the user who wishes to refer to the set B and employs a term $A \in V(S)$ lists the elements in B which are excluded from A or appended to A and bears a “cost” measured by the number of elements which have to be excluded or appended (for example, the sentence “You may eat only bread or any fruit with the exception of apples and bananas” utilizes the term fruit and three individual names: bread, apples, and bananas).

For a set B and a vocabulary V , define $\delta(B, V) = \min_{A \in V} d(A, B)$, the distance of the set B from the closest set in the vocabulary V . By assigning equal “probabilities” *a priori* to all possible sets that the user might wish to refer to, the problem for the designer becomes $\min_S \sum_{Z \subseteq X} \delta(Z, V(S)) / 2^{|X|}$. Essentially, the problem boils down to the choice of the *number* of elements in the optimal set S .

To *further* clarify the nature of the problem, we will carry out a detailed calculation for a four-element set X . Consider the case in which S contains two elements. The user can refer to any of the sets X , $\emptyset$, S or $-S$, without incurring any loss. He can approximate any one- or three-element set by using the set $\emptyset$ or the set X , respectively, while bearing a “cost” of 1. If he wishes to refer to

Choosing the semantic properties of language

one of the four two-element sets which are not S or $-S$, he incurs a loss of 2. Thus, the average loss is $[4(0) + 8(1) + 4(2)]/16 = 1$. A similar calculation leads to the conclusion that a choice of S as $\emptyset$ or X leads to an expected loss of $5/4$ (no loss for $\emptyset$ and X , a loss of 1 for the eight sets of size 1 and 3, and a loss of 2 for the six sets of size 2 – i.e. $[2(0) + 8(1) + 6(2)]/16 = 5/4$). If S is taken to be a one-element set (or a three-element set), the average loss is only $[4(0) + 6(1) + 6(1)]/16 = 3/4$. Thus, splitting X into two *unequal* subsets of sizes 1 and 3 minimizes imprecision.

Despite the above results, our intuition is correct in that the optimal size of S should be half that of X . When the set X is “large,” the loss associated with choosing an S containing half of X ’s members is “close” to being minimal. The following proposition is an exact statement of this result (its proof, omitted here, includes a combinatorial calculation):

Claim: Let X be an n -element set. The difference between the solution of the problem $\min_{S \sum_{Z \subseteq X} \delta(Z, V\{S\})/2^{|X|}}$ and the expected loss from the optimal use of the vocabulary spanned by an $(n/2)$ -member subset of X , is in the magnitude of $1/n^{1/2}$.

1.4 Informativeness

We now return to the world of binary relations on some set Ω . An additional function of binary relations on a set Ω is to transfer or store information concerning a specific relationship existing between the elements of Ω . Consider the case in which the grand set includes all authors of articles in some field of research and the speaker is interested in describing the relation “ x quotes y in his article.” The speaker may describe the relation by listing the pairs of authors who satisfy the relation. Alternatively, he may use those binary relations which are available in his vocabulary to describe the “ x quotes y ” relation. If he finds his vocabulary insufficient to describe the relation, he will use a binary relation which is the best approximation. For example, if the relation “ x is older than y ” is well defined,

the speaker can use the sentence: "An author quotes another if he is older than himself." As this may not be entirely correct, he can add a qualifying statement such as "the exceptions are a who did not quote b (though b is older), and c who did quote d (though d is not older)." Such qualifying statements are the "loss" incurred from the use of an imprecise relation in order to approximate the "who quotes whom" relation.

Our discussion envisages an imaginary "planner" who is able to design only one binary relation during the "initial stage of the world." Of course, real-life language includes numerous relations and the effectiveness of each depends on the entire fabric of the language. The assumption that the designer is planning only one binary relation is made solely for analytical convenience.

The design of one binary relation allows the speaker to select one of four binary relations. For instance, he can state: "Every author is quoted by all others younger than him," or "Every author is quoted by all others not younger than him," and of course he can also state "Everyone quotes everyone" and "No one quotes anyone," which do not require familiarity with any binary relations. Given a relation R , we will refer to these four relations as the vocabulary spanned by R and denote it by $V(R)$. (Note that in defining the vocabulary spanned by R , we ignore other possibilities for defining a binary relation using R , such as statements of the type " x S y if there is a z such that xRz and zRy .")

We assume that the speaker who wishes to refer to a binary relation S will use a relation in $V(R)$ which is the "closest" approximation. The loss is measured by the number of differences between the relation which the speaker wishes to describe and the one he finds available in his vocabulary. The distance between any two binary relations R' and R'' is taken to be the number of pairs (a, b) for which it is not true that $aR'b$ if and only if $aR''b$. Note that according to this measure, any pair for which R' and R'' disagree receives the same weight. Regarding the initial state, it seems proper to put equal weights on all possible "imprecisions." The designer's problem is to minimize the expected loss resulting from the optimal use of his vocabu-

Choosing the semantic properties of language

lary. It is assumed that from his point of view, all possible binary relations are equally likely to be required by the speaker. Thus, the designer's problem is $\min_R \sum_S \delta(S, V(R))$, where $\delta(S, V(R)) = \min_{T \in V(R)} \delta(S, T)$ and $\delta(S, T)$ is the distance between the relations S and T .

Stated in this way, the problem becomes a special case of the problem discussed in the previous section. The set X , in the terminology employed in that section, is the set $\Omega \times \Omega - \{(\omega, \omega) | \omega \in \Omega\}$. A binary relation R is identified with the graph of R – i.e. the subset of X of all (a, b) for which aRb . Recall that we concluded earlier that splitting the set X into two equal sets is nearly optimal for the designer if he wishes to reduce the expected number of “imprecisions.”

We have one more step to go to reach the goal of this section. In planning a binary relation on Ω , the designer may also consider the possibility that the relation will eventually be used in reference to a subset of Ω . This is analogous to a binary relation being indication-friendly if it allows the indication of any object out of *any subset of objects* (see section 1.2). Hence, R has to be “optimal” for potential use in indicating a subset of $\Omega' \times \Omega' - \{(\omega, \omega) | \omega \in \Omega'\}$ for every subset $\Omega' \subseteq \Omega$. If R is complete and asymmetric, for every subset $\Omega' \subseteq \Omega$, the induced relation $R_{|\Omega'}$ (defined by $aR_{|\Omega'}b$ iff $a, b \in \Omega'$ and aRb) includes exactly half the pairs in $\Omega' \times \Omega' - \{(\omega, \omega) | \omega \in \Omega'\}$.

In summary, our fictitious planner wishes to design a binary relation which spans a vocabulary with the goal of minimizing the expected inaccuracy of the term the user will actually use. Viewing a relation as a set of pairs of elements, the problem was linked to the optimization problem discussed in the previous section. There we concluded that splitting a set into two subsets allows “close to optimal” use of the induced vocabulary. Requiring the relation to be complete and asymmetric guarantees that for any subset of the grand set, the restricted relation will be “close to optimal” as an aid to the user in specifying a relation on the subset.

Conclusion 2: In order to express binary relations as accurately as possible on any subset of a set Ω using a vocabulary spanned by a single binary relation on the set

Ω , a binary relation on Ω which is complete and asymmetric is close to optimal.

1.5 Ease of desribability

In this section we discuss the third and last criterion by which binary relations are assessed in this chapter. Imagine that a hunter wishes to instruct his son on how to behave in the forest. When he observes two potential animals a and b , should he pursue a or b ? The instructions have to be applicable to any pair of animals and must be clear. Thus, the set of instructions can be represented by a complete and asymmetric binary relation R where aRb means that when the son simultaneously observes a and b he should pursue a .

The son is aware of the sets of animals that are edible and the structure of the relation R (i.e., the list of properties satisfied). The son is acquainted with the structure of the relation either because it is instinctual or because his father has informed him of these properties. Therefore, all that is left for the father to do, when transferring the content of R to his son, is to provide him with a list of "examples" – i.e., statements of the type "animal a should be pursued when you see it together with animal b ." The examples should be rich enough to allow the son to infer the entire relation from the structure and examples.

To illustrate, consider the case in which the relation R is a linear ordering and the number of elements in Ω is n . The minimal number of examples that the father must provide in this case is $n - 1$ ($a_1Ra_2, a_2Ra_3 \dots, a_{n-1}Ra_n$).

This brings us to the main topic of our discussion. We assume that providing examples is costly. (The complexities of the structure and process of making the inferences are ignored here.) The following problem comes to mind: What are the structures of the complete and asymmetric binary relations (tournaments) for which the number of examples required for their description is minimal?

Definitions: We say that $(f, \{a_iRb_i\}_{i \in I})$ defines the binary relation R^* on Ω when

Choosing the semantic properties of language

- f is a sentence in the language of the calculus of predicates with one binary relation named $a_i, b_i R$ and for all $i, a_i, b_i \in \Omega$
- R^* is the unique binary relation on Ω satisfying the sentence f and for all i , it is true that $a_i R^* b_i$.

The *complexity* of R^* , denoted $\ell(R^*)$, is the minimum number of examples that needs to be appended to R^* in order to enable a definition of R^* – i.e., the minimum size of the set I on all possible definitions of R^* .

As previously stated, our attention is limited to binary relations which are tournaments. The mathematical problem we wish to solve is:

$$\min_{R^* \text{ is a tournament}} \ell(R^*).$$

Note that for a given (finite) set of objects $\Omega = \{a_1, \dots, a_n\}$, the “structure” of any binary relation R^* can be expressed by the sentence $\varphi_{R^*}(x_1, \dots, x_n) = \exists x_1, \dots, x_n (\bigwedge \{x_i R x_j \mid a_i R^* a_j\})$. Thus, the optimization problem proposed above is equivalent to the following “puzzle-type” problem: Start with the graph of a tournament in which the names of the vertices have been erased. What is the minimum number of examples required to recover the names of the vertices (up to isomorphism)?

Example 1: Consider the tournament R^* on the set $\Omega = \{a, b, c\}$ where $a R^* b$, $b R^* c$ and $c R^* a$. The sentence which states that R is complete, asymmetric, and anti-transitive ($\forall x, y, z (x R y \text{ and } y R z \rightarrow \neg x R z)$) is consistent with two relations on Ω ; hence, a single observation, $a R b$, is needed to complete the definition of R^* . Thus, $\ell(R^*) = 1$.

Example 2: Let R^* be a linear relation. The relation is defined by a sentence f expressing completeness, asymmetry and transitivity, and by $n - 1$ examples $\{a_i R^* a_{i+1}\}_{i=1, \dots, n-1}$, where $a_1 R^* a_2 R^* a_3, \dots, a_{n-1} R^* a_n$. Obviously, there is no definition of a linear relation with less than $n - 1$ observations. Thus, $\ell(R^*) = n - 1$.

Example 3: Let $\Omega = \{a, b, c, d\}$ ($n = 4$) and let R^* be the relation satisfying $a R x$ for all x and $b R c R d R b$. R^* is defined by

$\exists wxyz[wRx \wedge wRy \wedge wRz \wedge xRy \wedge yRz \wedge zRx]$ and the three observations aRb , aRc and bRc .

Example 4: Consider the relation R^* on the set $\Omega = \{a, b, c, d\}$ described in figure 1.2. The structure of the relation R can be formulated by a sentence of the type $\exists v_1 v_2 v_3 v_4 \varphi(v_1, v_2, v_3, v_4)$. Twenty-four different binary relations on Ω have this structure. It is simple to verify that $\ell(R^*) = 4$.

Example 5 (Fishburn, Kim and Tetali, 1994):




Figure 1.2

This relation, R^* , defined on the 5-element set $\Omega = \{a, b, c, d, e\}$ satisfies $\ell(R^*) = 3$. It is nicely defined by a sentence expressing the property that for every x there are precisely two elements “beaten” by x . The three observations aRb , aRc , and eRb , define the relation through the chain of conclusions $\{dRa, eRa\}$, $\{cRe, dRe\}$, $\{bRd, cRd\}$ and, finally, bRc .

Examples 2 and 5 illustrate that for $n = 3$ and $n = 5$, a linear ordering is not the most “economical” structure. Are there any other binary relations with $n > 7$ which are defined by less than $n - 1$ observations? I am not aware of a complete answer to this question. However, we do know the following (this proposition was suggested to me by Noga Alon):

Proposition 1.3: For any ϵ there exists $n(\epsilon)$ such that for any $n > n(\epsilon)$ and for any complete and asymmetric relation R on a set of n elements, $\ell(R) > (1 - \epsilon)n$.

Choosing the semantic properties of language

Thus, at least for large sets, linear orderings are “almost” optimal with respect to the criterion of minimizing the number of observations required for their definition.

Comment: Notice that in the above discussion, we allowed the relation to be defined by a formula which depends on the number of elements in the set Ω . In contrast, the properties of linear orderings are expressed by a formula which does not depend on the number of elements in the set Ω . This leads to the following conjecture:

Conjecture: Let φ be a sentence in the language of the calculus of predicates which includes a single name of a binary relation R . There exists n^* such that if $|\Omega| \geq n^*$; then, for any tournament R^* which is defined on the set Ω by the sentence φ , $\ell(R^*) \geq |\Omega| - 1$.

In other words, although one can define a relation for “small” sets with less than $|\Omega| - 1$ examples, it is conjectured that a sentence must be accompanied by at least $|\Omega| - 1$ elements in order to define a tournament when the size of the set is “large enough.”

Comment: We conclude this section with an explanation of why the term “describability” – and not “learnability” – is used in this section. In our scenario the father chooses the examples he presents to his son. The number $\ell(R^*)$ is the minimum number of examples the father has to present in order to convey enough information for the son to deduce the content of R^* . When choosing the examples, the father knows the relation which he would like his son to learn. On the other hand, had the son wished to acquire the content of the relation R^* by asking a list of “questions” of the type “what should be chased, a or b ?,” he would not necessarily have asked first for the $\ell(R^*)$ “examples” which could convey the relation R^* . For example, he would need 2.5 inquiries “on average” in order to infer a linear ordering defined on a three-element set.

1.6 Discussion

This chapter investigated the observation that in a natural language, certain structures of binary relations appear much more frequently than others; in particular, we discussed properties that are satisfied by linear orderings. It was argued that certain functions of binary relations in natural languages are better served by relations satisfying these properties. I believe that this is no more than an interesting fact. A stronger interpretation of the results requires establishing a connection between the optimality considerations presented above and the realization of the optimal solutions in the real world. Such a connection would require at least three premises:

- (1) In order to function, natural languages include only a small number of structured binary relations.
- (2) Binary relations fulfill several functions in natural languages.
- (3) There are forces (evolution or a planner) which make it more likely that structures which are “optimal” with regard to the functions of binary relations will be observed in natural languages.

The first premise states that language inherently exhibits few of the properties of binary relations. Only if the number of possible structures is small can a user of the language deduce a relation’s structure from a small number of instances in which the relation is used. (It is amazing how few observations of the type “*a* is better than *b*” are sufficient to teach a child that this relation is transitive.)

The second premise is the central one in this discussion. There are numerous potential criteria by which to measure the functionality of binary relations. Three such criteria were examined above. It was argued that certain properties (all shared by linear orderings) perform better according to each of these criteria.

The third premise, which links the first two, states that either there is a linguistic “engineer” who chooses the properties of binary relations so that they function effectively or that evolutionary forces select structures which

Choosing the semantic properties of language

are optimal or nearly so with respect to the functions they fulfill. This idea, which is popular in economics, has also been noted by philosophers. For example, Quine states: "If people's innate spacing of qualities is a gene-linked trait, then the spacing that has made for the most successful inductions will have tended to predominate through natural selection" (Quine, 1969, p.126).

The approach adopted in this chapter is related to the functionality of the language approach discussed in linguistics (see, for example, Piatelli-Palmarini, 1970) and to attempts to explain the classification system in natural language (see Rosch and Lloyd, 1978).

The discussion in this chapter is also related to classical philosophical discussions on "natural kinds." The notion of a "natural kind" emerges from the philosophical inquiry into the factors which confirm an inductive argument (see, for example, Goodman, 1972, Quine, 1969, Watanabe, 1969). A key puzzle in this literature is the so called "riddle of induction": Let us say that up to this moment, all observed emeralds were green. Why does this observation imply that all emeralds are green rather than all emeralds are "grue," an alternative category which includes all objects which were green up to this moment and blue from now on?

One possible answer to this question is that the inductive process relies on notions of similarity. Inductive arguments are made only with regard to categories of similar objects. The category green contains similar elements; grue does not. A green element yesterday and a green element tomorrow are similar; in contrast, a grue element yesterday is not similar to a grue element tomorrow. However, this only begs the question since one is left with the problem of determining the natural similarity relations. Here we run into similar difficulties. One possible solution is to argue that two objects are similar if "most" unary predicates coincide in satisfying the two objects. The difficulty with this criterion is raised by "The Ugly Duckling Theorem" (see Watanabe, 1969, Section 7.6). If the set of predicates is closed under Boolean operations, then the number of predicates which satisfies any possible

object is constant; thus, the existence of elementary unary predicates cannot be the basis for explaining the existence of specific similarity relations. This led Watanabe to conclude that "if we acknowledge the empirical existence of classes of similar objects, it means that we are attaching nonuniform importance to various predicates, and that this weighting has an extra-logical origin" (Watanabe, 1969, p. 376). Thus, there is no escape from assuming that a certain kind of predicate (like "green" and not like "grue") has a preferred status called "natural kind" (see Quine, 1969).

Within the class of properties of binary relations, linear orderings, more than other structures, appear to be of a "natural kind." This chapter has attempted to provide some rationale as to why this is so.

REFERENCES

- Fishburn, P., J.H. Kim and P. Tetali (1994) "Tournament Certificates," AT&T Bell Labs, mimeo
- Goodman, N. (1972) *Problems and Projects*, Indianapolis and New York: Bobbs-Merrill
- Piattelli-Palmarini, M. (ed.) (1970) *Language and Learning: The Debate between Jean Piaget and Noam Chomsky*, Cambridge, MA: Harvard University Press
- Quine, W.V. (1969), *Ontological Relativity and Other Essays*, New York: Columbia University Press
- Rosch, E. and B. Lloyd (eds.) (1978) *Cognition and Categorization*, Hillsdale: Lawrence Erlbaum Associates
- Rubinstein, A. (1996) "Why are Certain Properties of Binary Relations Relatively More Common in Natural Language?," *Econometrica*, 64, 343-56
- Watanabe, S. (1969), *Knowing and Guessing*, New York: John Wiley