

LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

-72

Dewey
MASS. INST. TECH
MAY 11 1972
LIBRARY

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

THE PRACTICAL JUSTIFIABILITY OF
"IRRATIONAL" AVERSION TO RISK*

W. Wilcox
Jarrod W. Wilcox

March 1972

597-72

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139



MASS. INST. TECH.
MAY 9 1972
DEWEY LIBRARY

THE PRACTICAL JUSTIFIABILITY OF
"IRRATIONAL" AVERSION TO RISK*

11/11/72
Jarrod W. Wilcox

March 1972

597-72

*Presented at the International Meeting of TMS in Houston, Texas,
April 1972.

HO28
.M414
no. 597-72

RECEIVED
MAY 11 1972
LIBRARIAN

ABSTRACT

It is often necessary to explicitly confront cognitive limitations in prescribing appropriate decision-making behavior. One aspect of this general dictum is illustrated by showing that cognitively limited decision-makers may justifiably express greater risk-aversion in solving local problems than that allowed by current utility axioms. This occurs, for example, where strong externalities do not allow a satisfactory formulation of the event space. Consequently, the imposition of the expected utility hypothesis may sometimes be counter-productive. It is recommended that specific costs for risk be used to patch-up the "holes" in the choice space created by such externalities.

ACKNOWLEDGEMENTS

Without implicating them for the result, I would like to thank Gordon M. Kaufman, Paul A. Samuelson, Paul R. Kleindorfer and Murat R. Sertel for their comments.

I. INTRODUCTION

My aim is to persuade management scientists, and social scientists generally, of a specific need for modification of decision theory as a prescriptive guide to real-world decision-making. The issue raised here was raised and apparently resolved in the early nineteen-fifties. It reappears when one recognizes the necessity of explicitly confronting cognitive limitations in a prescriptive theory of management.

Management science relies on the axiomatizations of microeconomics and statistical decision theory for a normative view of good decisions. These generally incorporate the expected utility hypothesis.¹ By allowing subjective utility to be a non-linear function of objective pay-offs, the Von Neumann-Morgenstern (N-M) and the Savage axioms can be interpreted as allowing a degree of risk aversion. My thesis is that in a world of limited cognitive capability decision makers may justifiably express greater risk aversion than is consistent with these axioms. Unless we recognize these limitations, the empirical application of our normative approach may fail to be of real service.

Decision-Theory

The primary source of present axiomatizations of the theory of decision-making is the treatment given by Von Neumann and Morgenstern (1947). The specific axioms they note as required for utility functions are as follows.

Consider a system U of entities u, v, w, \dots . In U a relation $u > v$ is given, along with an operation $\alpha u + (1-\alpha)v = w$, where $\alpha, \beta, \gamma, \dots$ are real numbers greater than zero and less than unity. These concepts are

¹That is, that the utility of a lottery $\alpha u + (1-\alpha)v$ is equal to α times the utility of u plus $(1-\alpha)$ times the utility of v .

presumed to satisfy the following axioms:

Set Theoretic:

(3:A) $u > v$ is a complete ordering of U

Topological:

(3:B:a) $u < v$ implies $u < \alpha u + (1-\alpha)v$

(3:B:b) $u > v$ implies $u > \alpha u + (1-\alpha)v$

(3:B:c) $u < w < v$ implies there exists an α such that $\alpha u + (1-\alpha)v < w$

(3:B:d) $u > w > v$ implies there exists an α such that $\alpha u + (1-\alpha)v > w$

Algebraic:

(3:C:a) $\alpha u + (1-\alpha)v = (1-\alpha)v + \alpha u$

(3:C:d) $[\alpha(\beta u + (1-\beta)v) + (1-\alpha)v] = [\gamma u + (1-\gamma)v]$, where $\gamma = \alpha\beta$.²

The entities may be interpreted as events and their combining coefficients as probabilities. Savage (1954) extended this interpretation to include subjective probabilities.

The general acceptance of the usefulness of these axioms has not been without exception. In the early 1950's the so-called Archimedean axioms (3:B:c) and (3:B:d) were challenged.³ There have been many descriptive challenges of (3:A), the complete ordering axiom, and a few normative objections.⁴ Von Neumann and Morgenstern (N-M) themselves noted the inability of their system to adequately treat a specific utility or disutility for gambling. They suggested that alterations might have to be made in (3:C:d), the axiom concerning the algebra of combining lotteries of events,

²The equation numbering is from Von Neumann and Morgenstern. Of course the axioms of logic and basic set theory may also be presumed.

³See Thrall (1954).

⁴For example, H. Simon's suggestion of the real-world usefulness of satisficing represents such an objection. See Simon (1955).

in order to do so.

Although there are a number of more recent axiomatic treatments of decision theory, particularly in terms of mixture spaces, none of these give up Axioms (3:B:a) and (3:B:b), the so-called axioms of monotonicity. In the next section a class of occasions will be illustrated where monotonicity of preferences for lottery mixtures of events as perceived within cognitively limited frameworks is of doubtful justifiability.

II. A REAL WORLD SITUATION

In the real world, cognitive limitations are inescapable for any decision-maker. A valid prescription for the behavior of the decision-maker may in some instances have to confront and deal with these cognitive limitations. Let us informally describe an example.

Suppose that an individual decision-maker is faced with the possibility of two events, X and Y. He may take a variety of actions (or follow a variety of strategies), each of which may be expected to result in some lottery $\alpha X + (1-\alpha)Y$, where $0 \leq \alpha \leq 1$. Based on local information, he prefers $(1)(X) + (0)Y$ to $(0)X + (1)Y$. What can we say about his preference for intermediate lotteries? Is the answer obvious?

Suppose that the decision-maker is the manager of the production department in a business organization. Event X represents having produced 1000 mini-computers for inventory in January. Event Y represents having produced 100 mini-computers for inventory. Suppose this production manager prefers X to Y because it results in favorable manufacturing efficiency variances for the production department. Suppose there exists a sales department, which similarly confronts events R and S, where R represents

sales of 1000 and S represents sales of 100. R is preferred to S within the sales department because it results in higher sales commissions.

Suppose the ultimate preferences of the managers of both departments are determined not only by their preferences based on local conditions but also by rewards from top management which depend on the global performance of the business organization.

Suppose that one may rank order the global performance of the organization according to the ordered pairs of certain (ex post) events as follows:

(X,R)

(Y,S)

(X,S)

(Y,R).

That is, the local outcomes in the production and sales departments are complementary in producing global outcomes. One can think of this situation as the confluence of two branches in a "means-ends tree".⁵ The local "ends" are complementary "means" to a global "end". Suppose these decisions are made at the beginning of the month, and that ex post global performance affects rewards to the production manager only at the end of the month.

An aspect of the real world which is of crucial importance for our example concerns the costs of communication between cognitively-limited decision-makers. If we assume these costs to be sufficiently low, we avoid the issue of cognitive limitations altogether by permitting aggregations of intelligence. Therefore, suppose that communication between the production

⁵See Newell and Simon (1971).

and sales managers is limited by various costs, so that we must confront a limited information structure.

Now, one could presume any of a fairly rich variety of limited information structures. Each would reflect different types of cognitive limitations. The following are a non-exhaustive list.

Condition A:

For example, the production manager could conceivably be unaware of the nature of the sales department and global performance, but learn through repeated monthly trials that those of his decisions associated with risky lotteries of X and Y tended to be followed by decreased rewards from top

management.⁶

Condition B:

He may, on the other hand, be aware of the sales department but unable to enumerate its outcome set, and thus clearly unable to formulate the global problem. Because of both practical (cognitive limitations) and theoretical (there is no finite stopping rule) reasons neither the production nor the sales manager views the sequence of monthly decisions as a single problem, but rather as a series of distinct monthly problems whose character may change through time. The production manager may, however, know that the sales manager is similarly rewarded by top management based on global performance, and that global performance depends in a complementary way on production and sales outcomes. Further, he may know that past production outcomes have been very predominantly Y for a long enough sequence for the sales manager's decisions to have come into complementary equilibrium with Y through the mechanisms of Condition A.

Condition C:

The production manager may be in condition B, with the exceptions that there is no history of stable past outcomes and that the sales manager is perceived as able to formulate and solve the reduced global problem incorporating a fixed lottery of X and Y. That is, the sales manager is assumed to be able to coordinate his activities with production on the basis of communication of one period's production outcome lottery. Note, however, that even the sales manager is still not assumed to be able to formulate the full

⁶See Estes (1954),

global problem.

Conditions A, B, and C each lead to the following argument.

Let us consider the production manager's justifiable preferences for various outcome lotteries of X and Y and thus for various activities which will produce these lotteries. Suppose his reward from top management varies positively and monotonically with global performance. Lotteries where X and Y are both likely will increase the probability of global outcomes (X,S) and (Y,R), decreasing his rewards from top management. Thus, the production manager will be justifiably risk averse. Possibly the trade-off between changes in local rewards and changes in the rewards from top management will be such that his degree of risk averseness may be accommodated within the framework of monotonicity. In such a situation, represented in Figure I, his risk-averseness may be represented by a curved but always increasing utility curve for certainty lotteries.

In general, however, the differences in global performance between outcome pairs (X,R) or (Y,S) and (X,S) or (Y,R) may be arbitrarily large. Consider Condition A. The change in reward from top management stemming from global performance will, beyond some point, outweigh the purely local considerations of favorable production efficiency. Consider Conditions B and C. Suppose the production manager believes that the complementarities, without knowing their nature, between production and sales are sufficiently strong so that any substantial probability of uncoordinated behavior will produce negative rewards from top management which will outweigh local preferences.

We see that in each of these conditions, A, B, or C, the production manager may justifiably be more risk averse than is consistent with the

Subjective
Utility of
Lotteries




Figure I

"Rational" Risk Aversion

N-M axioms, without being able to formulate the global problem. That is, the production manager, even though he prefers $(1)X+(0)Y$ to $(0)X+(1)Y$, will prefer both to $(\alpha)X+(1-\alpha)Y$ for some α . Thus the management scientist's attempt to empirically construct a utility function for lotteries of X and Y will fail, and the alternative of enlarging the problem will not necessarily be possible.

III. DISCUSSION

The well-trained management or social scientist will, of course, immediately raise an objection to this example. He will note that the difficulty lies in misspecification of the outcome event space. Thus, he will disavow any threat to the expected utility hypothesis. In such situations, the proper application of statistical decision theory requires taking the global viewpoint. But, comes our reply, what if the decision-maker, because of cognitive limitations, as in Conditions A, B, or C, cannot take a global viewpoint? Well, says the scientists, statistical decision theory is normative, not descriptive. Again, we reply, what should the decision-maker do? Perhaps, he weakly says, the theory is not always prescriptive. And therein lies my thesis, which parenthetically, is not original. Savage (1954) noted the problem at some length; also, the presumption of the absence of externalities was discussed in passing by Samuelson (1952).⁷

⁷See Savage (1954, pp. 82-88). Samuelson (1952, p. 677) also raises the problem: "In what dimensional space are we 'really' operating? If every time you find my axiom falsified, I tell you to go to a space of still higher dimensions, you can legitimately regard my theories as irrefutable and meaningless." But then he goes on to say that the necessary freedom from externalities may be a reasonable approximation in most cases of interest to economics. See also Wold (1952) and Samuelson (1952) for an early recognition that time sequences of local choices may represent global problems of a similar nature. My arguments with respect to "cross-sectional" local versus global difficulties may all be transposed into a discussion of individual decisions in stochastic processes.

Unfortunately, the absence of externalities as a restriction on the N-M or Savage axioms' applicability seems to have been little noticed by those who have followed. The possibility of externalities interfering with attempts to empirically measure subjective utilities and subjective probabilities has been largely overlooked.⁸ Apparently, only Savage noted the problem of practically locating appropriate microcosms within which the axioms would hold--that is, those with no strongly interfering externalities, or those with "strong independence."

Work in Utilities for Multi-dimensional Attributes

There exists a considerable body of theoretical, and some empirical, work which might be confused with our topic. This existing body of knowledge deals with evaluating multi-dimensional consequences under uncertainty.⁹ Its major thrust is to show under what conditions utilities for single attribute choices can be combined to give utilities for multi-attributed choices. The conditions first described included a restriction to situations where global utility depends only on marginal probability distributions over each attribute and not on joint probability distributions over more than one attribute. That is, there is no complementarity. However, in recent years, a theory of additively separable utility functions has been developed which permits the empirical evaluation of global utility from individual

⁸See Grayson (1960), Raiffa (1961), Davidson, Suppes, and Siegel (1957), Mosteller and Nogee (1951).

⁹See Keeney (1969), Fishburn (1964, 1967), Gorman (1959). This theory can also be applied to certain sequential as well as cross-sectional problems.

attribute utilities in certain cases where there is a simple form of complementarity. Keeney (1970) has noted that many of the difficulties encountered empirically in attempting to apply "decision analysis" stem from problems in trying to construct what I would call global utility functions. However, he takes a very different approach than the present. In our production and sales examples, he would show under what conditions local utility functions in the production and sales departments might be combined to produce a global utility function. On the other hand, I have not assumed the existence of either local or global utility functions, and am in fact specifically attacking the empirical and axiomatic basis under which one is led to assume the existence of local utility functions.

Such work as that by Keeney may be useful in gaining some insight into when our problem will be severe. Clearly, under conditions of additively separable utility and the existence of a global utility function, we now know that local utility functions are possible. However, in a practical sense we have not come very far in determining, in Savage's terms, appropriate microcosms. Still less do we know what to do when we are stuck with an inappropriate one.

The Generality of the Potential Difficulty

Complementarities are characteristic wherever two distinct means are used jointly to obtain a given end. Decentralization of some decision-making activity to separate, cognitively-limited decision-makers governing these means is also omnipresent, again because of cognitive limitations. One is hard put to think of any outcome to a real world decision that could not be viewed as a partial means to a more global end. Thus, the existence

of unavoidable externalities in applying local decision-theoretic analysis is extremely widespread, if not universal. For us as management scientists, a key empirical issue in applying decision analysis is determining whether these externalities are sufficiently strong to make unjustifiable the Savage axioms' application. This is an unsolved problem. However, general considerations suggest that we look for problems where a high degree of cognitive capability can be brought to bear on a small problem domain buffered in some way from strong externality feedbacks.

IV. ANOTHER SITUATION OF SPECIAL INTEREST

Game theory is based on the N-M axioms regarding preferences on outcomes. Suppose we regard the outcome of a game in terms of $(A,0)$ pairs, where A is the action or strategy followed and 0 is the conventionally defined payoff lottery given mini-max conditions. Let us regard the decision-maker's preferences for the 0 's as satisfying the N-M axioms, including monotonicity. What conditions will the decision-maker's preferences for various A 's satisfy? We know game theory shows that for many games the decision-maker will normatively prefer a mixed strategy. That is, he will prefer a risky lottery of strategies to his most preferred pure strategy. Clearly, then, his preferences for the A 's will not satisfy the monotonicity condition imposed on his preferences for the 0 's. One could even say he had a utility for gambling. Specifically, N-M Axiom (3:B:b) is violated for the A 's, though not for the $(A,0)$ pairs.

Are there real-world cases where the decision-maker can formulate the problem in terms of the A 's, but not in terms of the $(A,0)$ pairs? Suppose

that the analysis of the game structure which led to preferences for $(A,0)$ pairs and thus for A 's is forgotten, but that the result, that is, a preference ordering over the A 's, is remembered. Or suppose one individual analyzes the game and tells another what to do. Then the result is obvious. Some, of course, would object to the use of strategies as events over which one is presumed to have preferences. This objection, however, implicitly presumes that one can recover the outcome space and game structure for which strategies are means. In a world of cognitive limitations, this recovery is not always feasible.

The $(A,0)$ event space can be used to further illuminate our earlier production-sales case. Consider Condition A. Let us further suppose that the production manager can formulate his local problems in terms of $(A,0)$ pairs, where the A are lotteries of X and Y , and where the 0 are themselves ordered pairs (A,T) , where T are top management reward events. That is, we restrict ourselves to the case where the 0 can be effectively described, in the way that the global outcome space could not be. Finally, we need not assume, as in the earlier game example, that the association between the A and 0 has been forgotten.

First, suppose the manager can formulate the problem in terms of the $(A,0)$ or of the 0 , in accordance with the N-M or Savage axioms. By formalism, we have in this case indeed altered the problem to that of multi-dimensional utilities because 0 is multi-dimensional. Again, however, lest we jump to the conclusion that we have reached familiar ground, note that we have no assurance of the existence of a local utility function for the X,Y lotteries, (or even for the T). And clearly, nothing has changed with regard to the inapplicability of the

expected utility hypothesis to these X,Y lotteries.

Second, we have again the problem of worrying about cognitive limits in defining the problem. How can we be sure in practice that the (A,0) problem, now defined necessarily in an even more personal way, may not be local with respect to still more global personal problems of the production manager? Again, the argument is recursive.

V. IMPLICATIONS FOR THEORY

Currently, normative theory for individual decisions and microeconomics rest on the same base.

Existing non-deterministic microeconomics, including market and price theory, are built to some extent on a foundation of individual cardinal utilities for bundles of goods and services, which in turn rest indirectly on the Von Neumann-Morgenstern derivation of cardinal utility functions based on their axioms for preferences for lotteries. If one takes seriously the problem of externalities induced by necessary cognitive limitations the axiom of monotonicity or strong independence for individual decisions must be eliminated. In order to save non-deterministic microeconomics, one may repair these foundations at either of two points. The first of these could conceivably serve as a basis for weak, normative views on individual decisions.

Individual utility can be reconstructed based on the notion of axiomatizing the ability to directly compare differences (or similarities) in preferences. That is, suppose $X > Y$ and $X > Z$. If it is possible to rank order these pairs in terms of the degree to which the first event is preferred, and if other plausible axioms are presumed, cardinal utility functions can

be derived.¹⁰ For example, if $X \succ Y$ more strongly than $X \succ Z$, it is possible to deduce that the utility of Y is closer than that of Z to the utility of X, etc.

In the nineteen fifties, this axiomatization of cardinal utility was rejected in favor of the N-M approach. It was reasoned then that the comparison of strengths of preference was a less natural operation than the ability to express preferences for lotteries which satisfied the N-M axioms. However, as more recent experience in non-metric multi-dimensional scaling has shown, decision-makers experience little difficulty in ranking event pairs in terms of similarities of attributes. This should extend readily to ranking event pairs in terms of similarities of preferences and thus naturally operationalize the necessary axiom. In allowing multi-dimensional operationalization, there would also be benefits in dealing with multi-dimensional utilities, both of the well-behaved Keeney variety, and of the ill-behaved variety such as arise from lexicographic ordering.¹¹

Though convenient for the foundations of non-deterministic micro-economics, such a theory might not, however, be very useful as a normative theory for individual decision-making. First, much of the normative content of the Savage axioms would have been removed; second, the complexities of empirically measuring subjective utilities and probabilities under the new axiom set would be rather cumbersome.

¹⁰Though not if only adjacent differences are compared. See Samuelson (1947), Georgescu-Roegen (1954), and especially the later literature on the related issue of non-metric multi-dimensional scaling, for example, Shepard (1962) and Kruskal (1964).

¹¹See Thrall (1954).

The other possible point of fundamental reconstruction in non-deterministic microeconomics is in the development of the theory of markets without depending on cardinal utility functions for individuals. This makes sense in view of the role markets play in relaxing the constraints of individual cognitive limitations. The new theory would explain how many individual decisions made "irrationally" are aggregated into "rational" social decisions by the market mechanism. Of course I mean by "rational" merely a conformance to the existing theorems of microeconomics. Such a reconstruction would represent a great advance in theoretical sophistication. Thus, this task would not be so easy as patching up individual utility theory.

VI. IMPLICATIONS FOR POLICY ADVISEMENTS

Management scientists, along with economists, in particular, presuppose global theoretical frameworks on real-world policy which exists in microcosms. This happens not only in operations research but even in fairly broad economic cost-benefit studies. In all of these situations, externalities are present which are too complex to be fully incorporated in the analysis. When their effects are strong, they may invalidate the current basis for decision analysis by making unjustifiable the Savage axioms as applied to local problem formulations.

Sometimes the policy-maker or manager protests our approach, but he can't put his finger on the source of his intuition of unreality. Yet we, in our hubris, deride his "irrationality". And we wonder why our services are not sought for more vigorously.

The unpopular lesson that I would draw is this. Until we learn how to more effectively incorporate necessary cognitive limitations in our decision theory, we should slow in extending its norms and methods to the

realm of individual decision-making. Where market mechanisms relieve cognitive constraints, we may proceed with a little more confidence, but even there it would be inappropriate to close our eyes to further difficulties.

What can we, as management scientists, do until more progress is made toward effective theoretical recognition of cognitive limits is made?

First, one can concentrate on analyzing situations where large cognitive resources can be concentrated on a domain relatively free of important externalities. Second, when evaluating statistical decision trees where serious externalities are unavoidable, it may be wise to face up to our cognitive limits by incorporating a specific cost or benefit to risk in the analysis.

REFERENCES

- Estes, William K. "Individual Behavior in Uncertain Situations: An Interpretation in Terms of Statistical Association Theory" in Decision Processes, R. M. Thrall, C. H. Coombs, and R. L. Davis (ed.), John Wiley and Sons, Inc., New York, 1954.
- Fishburn, P. C. Decision and Value Theory, John Wiley and Sons, New York, 1964.
- Fishburn, P. C. "Interdependence and Additivity in Multivariate, Unidimensional Expected Utility Theory", International Economic Review, Vol. 8, pp. 335-342, 1967.
- Georgescu-Roengen, Nicholas. "Choice, Expectations and Measurability", Quarterly Journal of Economics, Vol. 68, pp. 503-534, 1954.
- Gorman, W. M. "Separable Utility and Aggregation", Econometrica, Vol. 27, pp. 269-481, 1959.
- Grayson, Charles Jackson. Decisions Under Uncertainty: Drilling Decisions by Oil and Gas Operators, Harvard University Division of Research, Graduate School of Business Administration, 1960.
- Keeney, R. L. "Multidimensional Utility Functions: Theory, Assessment and Application", Technical Report No. 43, Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1969.
- Keeney, R. L. "Assessment of Multiattribute Preferences", Science, Vol. 168, pp. 1491-1492, June 1970.
- Kruskal, J. "Multi-dimensional Scaling by Optimizing Goodness of Fit to a Non-Metric Hypothesis", Psychometrika, Vol. 29, pp. 115-129, 1964.
- Luce, Robert Duncan and Howard Raiffa. Games and Decisions; Introduction and Critical Survey, Wiley, New York, 1957.
- Newell, Allen and Herbert A. Simon. Human Problem Solving, Prentice-Hall, Englewood Cliffs, N.J., 1971.
- Raiffa, Howard. "Risk, Ambiguity and the Savage Axioms, Comment", Quarterly Journal of Economics, Vol. 75, pp. 690-694, 1961.
- Samuelson, Paul A. Chapter VII in Foundations of Economic Analysis, Harvard University Press, Cambridge, Massachusetts, 1947.
- Samuelson, Paul A. "Probability, Utility, and the Independence Axiom", Econometrica, Vol. 20, pp. 670-78, 1952.

Savage, Leonard J. The Foundations of Statistics, John Wiley and Sons, Inc., New York, 1954.

Simon, Herbert A. "A Behavioral Model of Rational Choice", Quarterly Journal of Economics, Vol. 69, pp. 99-118, 1955.

Thrall, Robert M. "Applications of Multidimensional Utility Theory", in Decision Processes, R. M. Thrall, C. H. Coombs, and R. L. Davis (ed.), John Wiley and Sons, Inc., New York, 1954.

von Neumann, John and Oskar Morgenstern. Theory of Games and Economic Behavior, (2nd edition), Princeton University Press, Princeton, N.J., 1947.

Wold, H. "Ordinal Preferences or Cardinal Utility?", Econometrica, Vol. 20, pp. 661-664, 1952.

Date Due

NOV 22 '76		
	MAY 22 1985	
MAY 5 '76		
APR 22 '77	OCT. 21 1976	
MAY 1977	FEB 2 1977	
NOV 29 '77		
MAY 16 '78		
SEP 13 '78		
MAR 17 '80		
MAR 31 '83		
		Lib-26-67

MIT LIBRARIES



588-72

3 9080 003 701 635

MIT LIBRARIES



581-72

3 9080 003 701 601

HD20

.M414

Nos. 5

590-72 Nos. 5

MIT LIBRARIES



3 9080 003 701 585

MIT LIBRARIES



592-72

3 9080 003 670 608

MIT LIBRARIES



593-72

3 9080 003 670 574

MIT LIBRARIES



594-72

3 9080 003 701 593

MIT LIBRARIES



591-72

3 9080 003 701 569

MIT LIBRARIES



595-72

3 9080 003 670 715

MIT LIBRARIES



596-72

3 9080 003 701 718

MIT LIBRARIES



597-72

3 9080 003 701 692

MIT LIBRARIES



598-72

3 9080 003 670 707

MIT LIBRARIES



599-72

3 9080 003 701 650

