DESCRIPTIVE AND PRESCRIPTIVE MODELS FOR JUDGMENT AND DECISION MAKING: IMPLICATIONS FOR KNOWLEDGE ENGINEERING

Thomas R. Stewart
National Center for Atmospheric Research
(Sponsored by the National Science Foundation)
Boulder, Colorado USA

Claude McMillan, Jr.
Center for Applied Artificial Intelligence
College of Business and Administration
University of Colorado at Boulder
Boulder, Colorado USA

INTRODUCTION

Development of a knowledge base for an expert system often requires an analysis of the inference processes and decision making strategies of one or more experts. Knowledge engineers therefore share the psychologist's and operations researcher's interest in inference and decision making. Psychologists have been interested in understanding and describing how people make inferences under complexity and uncertainty. Operations researchers have developed methods that extend human ability to make decisions about complex problems.

Both operations researchers and psychologists have developed methods for modeling inference and decision making. The psychologists' methods produce descriptive models: they describe how people actually behave. Operations researchers produce prescriptive models: they prescribe optimal strategies or efficient heuristics for coping with complex problems [1]. Researchers in both areas have produced models which have not only equaled expert performance, but have outperformed experts. Examples of research in these areas will be briefly described and implications for knowledge engineering discussed.

Stewart, T.R. and McMillan, C., Jr. (1987). Descriptive and prescriptive models of judgment and decision making: Implications for knowledge engineering. In J. Mumpower, L. Phillips, O. Renn, and V.R.R. Uppuluri (Eds.) Expert Judgment and Expert Systems. New York: Springer-Verlag.

DESCRIPTIVE MODELS OF INFERENCE

Knowledge engineers generally rely on experts' verbal descriptions of their reasoning. This approach (process tracing or protocol analysis) has been used in psychology as well (Kleinmuntz, 1968; Ericsson and Simon, 1984). Traditionally, however, psychologists have been suspicious of verbal reports, and they have developed methods for analyzing judgment that do not depend on a person's ability to provide accurate verbal descriptions of thinking processes.

One such method, called judgment analysis, derives the parameters of judgment from a statistical analysis of a sample of actual judgments (Hammond et al., 1975). The expert is required to do only what he or she does best, that is, to make judgments in a natural setting using familiar materials. The analyst then uses standard statistical methods to develop a model that describes the inference process that produced the sample of judgments. In other words, the analyst produces a statistical model of the expert's inference process.

The data required for the development of such a model are a number of cases of a particular type of judgment. Each case includes the information used to make the judgment and the resulting judgment. Cases may be obtained in a natural setting (e.g., loan application decisions obtained from bank files) or in controlled settings. In a controlled setting, an expert is asked to make judgments based on a sample of real or hypothetical scenarios. The size of the sample required is dictated by the standard statistical requirements for reliability of parameter estimates. Table 1 illustrates two scenarios that might be used in a study of judgments of bank loan officers. Wilsted et al. (1975) asked 165 bankers to evaluate 50 different loan applications, similar to those illustrated in Table 1.

The items of information that are available to the judge, called cues, are considered independent variables in a statistical analysis with the judgment as the dependent variable. Multiple regression analysis is typically used to analyze the data, but other methods may be more appropriate in particular applications. The analysis yields a model of the judge, which expresses the judgment as a mathematical function (generally a polynomial) of the cues, and an index of how well the model fits the judgments (e.g., the squared multiple correlation

Table 1. Hypothetical loan applications.

LOAN APPLICATION A

Amount requested: \$15,000
 Age: 36 No. of Dependents: 3
 Years with present employer: 4
 Monthly Income: 1950 Total loan payments: 758
 Average Checking Account Balance: 500
 Average Savings Account Balance: 3500
 Occupation: White collar Own or rent home: Own
 Credit rating (1-4): 4
 Marketability of Collateral (1-4): 3

LOAN APPLICATION B

Amount requested: \$12,500
 Age: 31 No. of Dependents: 3
 Years with present employer: 1
 Monthly Income: 1100 Total loan payments: 340
 Average Checking Account Balance: 300
 Average Savings Account Balance: 6000
 Occupation: Skilled Own or rent home: Rent
 Credit rating (1-4): 3
 Marketability of Collateral (1-4): 2

coefficient, or coefficient of determination). For example, in the Wilsted et al. study the mean of the squared multiple correlations for the 165 bankers was .709. Since we do not expect judges to be perfectly reliable, squared multiple correlations ranging from .7 to .9 are typical in judgment studies and indicate a reasonably good fit of the model to the judgments.

Figure 1 illustrates how a model based on a statistical analysis of judgments of one judge can be interpreted graphically for a simplified problem involving only three cues. The graphic interpretation of the regression model is based on transformations described by Hammond et al. (1975). The figure illustrates the decomposition of judgment into relative weights for each of the cues and the form of the functional relationship between each cue and the overall judgment. Such decomposition and graphic displays have produced useful insights into the judgment process.

Although the use of statistical models of inference can be found as early as 1923 (Wallace, 1923), extensive use of the method began in the mid-1950s (Hammond, 1955; Hoffman, 1960). Several computer programs for judgment analysis, including versions for personal computers, are available [2]. A description of Social Judgment Theory, which is the theoretical foundation for judgment analysis, and a description of the method itself can be found in Hammond et al. (1975). A variety of experts have been modeled including physicians (e.g., Fisch et al. 1981; Kirwan et al., 1983), stock brokers, clinical psychologists, and polygraph interpreters. Although the models were capable of approximating intelligent behavior in restricted domains, they were used only as research tools. No one thought of marketing them as expert systems.

Statistical Models vs. Human Judgment

One of the striking findings of the early research on statistical models of judgment was that, for both novice and expert judges, across a variety of problem domains, simple linear additive models provided a description of the judgment process that accurately reproduced the original judgments (Slovic and Lichtenstein, 1971). Even more striking was that the models consistently outperformed expert judges themselves in tasks where a correct answer was available (Goldberg, 1970). This occurred because the models excluded the residual component

RELATIVE WEIGHTS

FUNCTION FORMS

Figure 1. Graphic interpretation of a statistical model of judgment.

of judgment, which was unrelated to the criterion.

An important line of research comparing human judgment and statistical models involved diagnosis of mental illness based on the results of a psychological test called the Minnesota Multiphasic Personality Inventory (MMPI). The results of this paper and pencil test for an individual consist of scores on 11 scales. In the 1950s, use of the MMPI for diagnosis (e.g, of psychosis vs. neurosis) was thought to require an understanding of complex patterns of interactions among the 11 scales. Developers of the MMPI argued that without extensive training and experience clinicians could not make proper use of the scores, and they discounted the possibility that any simple model, such as a sum of weighted scale scores, could capture the richness of the diagnostic process (Goldberg, 1969). They felt that human judgment was an essential element of the process.

The clinicians' arguments for the complexity of the process of interpreting MMPI results and for the essential role of human judgment disintegrated as a result of a series of studies of clinical judgment:

1. When both experienced clinicians and graduate students judged whether patients were medically or psychiatrically hospitalized, the judgments of the experienced clinicians were no better than those of the graduate students (Oskamp, 1967).
2. Using psychiatric diagnoses (based on extensive interviews) of each patient as a criterion, statistical diagnoses of psychosis vs. neurosis, based on a multiple regression model of the relation between the MMPI scales and the correct diagnosis, outperformed the diagnoses made by the clinicians (Goldberg, 1969).
3. Diagnoses based on regression models of the clinicians' own judgments outperformed the clinicians themselves (Goldberg, 1970).
4. The clinicians were beaten by a model which simply added the 11 equally weighted scores (Dawes and Corrigan, 1976).

This work suggested a method called bootstrapping for improving judgment. Under certain conditions, accuracy can be

improved by replacing the expert with a statistical model of his or her own judgment process. Camerer (1981) cited 15 studies showing an improvement in accuracy due to bootstrapping. The studies included judgments of graduate school success, survival predictions based on cancer biopsies, diagnosis of mental illness, personnel interviewing, students' judgments of teacher effectiveness, and selection of life insurance salesmen. Armstrong (1985) reviews the evidence for the usefulness of bootstrapping. He cites examples of success involving judgments about electronics production and catalog sales, as well as forecasts of sales success of personnel, returns for securities, loan defaults, and bankruptcy.

Although the work on bootstrapping is intriguing and has yielded important insights into some of the limitations of expert judgment, it is not universally applicable nor does it guarantee improved accuracy. It can only be applied to problems which are amenable to statistical modeling (this point is taken up below), and the properties of task and judge that are necessary for the success of bootstrapping have been described by Camerer (1981).

Implications for Knowledge Engineering

It is interesting to speculate about what would have happened in the 1950s if an attempt had been made to use today's methods to develop an expert system for clinical diagnosis from the MMPI. Knowledge engineers would have had no reason to question the expert clinicians' views regarding the complexity and richness of the interpretation process. Extensive time and effort would have been spent trying to capture the subtleties of the diagnostic process. In fact, complicated objective rules for diagnosis from the MMPI were being developed in the 1950s (Meehl and Dahlstrom, 1960). These rules performed no better than simple linear models (Goldberg, 1969). Consequently, an expert system based on such rules would not have outperformed one developed in much less time using a simple linear regression analysis. The MMPI experience suggests that knowledge engineers should regard experts' descriptions of their inference processes with some skepticism and explore the utility of simple decision-making models.

Limitations of judgment analysis. We are not, of course, recommending that knowledge engineers discard verbal protocols

and production functions and adopt statistical models exclusively in their work. The research to date has necessarily involved static, rather than dynamic tasks (Hogarth, 1981) involving judgments based on a few (5-20) items of information, and generally there is both high uncertainty in the task and fairly high inconsistency in the judgments of the expert. Furthermore, application of the method is not always straightforward. Research has shown a number of limitations and potential pitfalls in the use of statistical analyses to develop models of judgment. Results not only depend on the context in which the judgments are generated, they also depend on how the data are analyzed. Although there are many situations where the modeling process can be more or less automatic, there are others where the expertise of an experienced judgment analyst is required. More research is needed to determine when such models might be useful in knowledge engineering.

The use of statistical methods to develop models of judgment has been criticized on the grounds that the inference process is so complex that statistical modeling is impossible or impractical. It is argued that, since judgments about important problems generally involve multiple variables that interact in complex patterns, statistical modeling would require an unreasonable number of cases, and would likely produce misleading results. There are undoubtedly problems for which this is the case. For example, no applications of judgment analysis to weather forecasting have been reported in the literature. (However, the results of Allen et al, 1986, indicate that simple statistical methods may prove useful in this context as well.) Experience has shown, however, that for many problems, statistical modeling is practical and useful. O'Brien (1985), for example, reports a successful application of judgment analysis to strategic decision making in R & D management and argues that judgment analysis can outperform knowledge engineering in this context. Even very complex problems can often be decomposed into subproblems, one or more of which may be amenable to judgment analysis.

Judgment analysis and verbal protocols. Judgment analysis offers a major advantage as a tool for knowledge engineering: It provides a method for modeling inference that does not rely on the expert's ability to describe his or her thinking process. This is important because judgmental competence is not always accompanied by the ability to describe accurately the process

that produced the judgments. Verbal descriptions of reasoning can be incomplete, inaccurate or misleading. Some important aspects of the expert's inference process may not be readily accessible and may be difficult to translate into words. The description obtained may be unduly influenced by the method used to elicit it. Questions posed by the knowledge engineer and the requirements of particular shells impose a frame on the problem. Seemingly irrelevant and inconsequential aspects of problem framing have been shown to have a powerful effect on judgment (Tversky and Kahneman, 1981). For these reasons, it is desirable to have a knowledge engineering tool that does not depend on the expert's ability to describe the inference process.

There is another reason that knowledge engineers should not rely completely on verbal descriptions. A series of descriptive studies of judgment, initiated by a classic paper by Tversky and Kahneman (1974), has identified a number of mental strategies that can lead to errors in judgments. Research results repeatedly recount systematic biases in judgment, the inconsistent and poorly controlled nature of the judgment process, the pervasiveness of cognitive limitations that can reduce the validity of judgments, and the difficulty of overcoming those limitations (Hogarth, 1980; Kahneman et al., 1982). Irrelevant features of judgment tasks can strongly influence judgment while relevant information is often ignored or used inappropriately. These biases and limitations have been found in both expert and lay judges. Although the generality of some of this research is currently being debated, and some researchers are adopting a more optimistic view of human cognitive ability (Ebbesen and Konecni, 1980; Nisbett et al., 1983; Kruglanski et al., 1984), the possibility that experts' cognitive processes are flawed cannot be discounted. For example, Stewart and Glantz (1985) described how the judgments of expert climatologists may have been biased.

Although verbal descriptions of reasoning are likely to remain a staple of knowledge engineering, they should not be the only tool. Judgment analysis can be used in combination with verbal protocols (Einhorn, et al., 1979) so that the knowledge engineer can take advantage of the insights provided by both while avoiding their pitfalls. If the knowledge engineer relies solely on the expert's verbal statements, then both the efficiency of system development and the ultimate performance of the system will be limited by the expert's ability to verbally describe his or her inference process.

PRESCRIPTIVE MODELS

Under the circumstances, one strategy for the knowledge engineer is to be aware of potential biases and limitations of expert judgment and to guard against them by using a combination of methods to elicit information about the inference process. But another strategy is to strive to improve the knowledge base through critical, in-depth thinking about the nature of the problem and by using analytic models that can, in many cases, far exceed the limited information processing capabilities of a human expert in certain problem-solving settings.

Since World War II, operations researchers have been approaching problems using such models, that is, by developing prescriptive models of decision making. By and large their efforts have been directed toward decision problems in which the important variables can be expressed in quantitative terms. Especially interesting to operations researchers have been management problems in resource allocation. Examples include:

- o determining how many of what products to produce in each of several alternative production facilities, and how the output should be transported to various possible distribution facilities in such a way as to minimize costs;
- o identifying that mix of investments which will maximize the return, commensurate with the investor's risk propensities;
- o assigning work to employees in such a way as to minimize the elapsed time required for completion of a project.

As the wording above implies, the emphasis has been not just on producing good decisions, but on producing optimal decisions, that is maximizing or minimizing with respect to some quantitatively specified criterion. Thus, operations researchers have sought to build knowledge bases that perform not only as well as, but in most cases significantly better than experts.

Managerial Decision Making

We predict that a substantial portion of the most useful work in knowledge based systems, in the years ahead, will focus on problems faced by managers in both the public and the private

sectors. The opportunities are substantial. Whereas managers themselves have tended to look at blue collar operations for opportunities for increased productivity, it is generally agreed among students of administration that the administrative hierarchy itself is the place to look for improvement opportunities. Of course, it is just such managerial problems that have occupied operations researchers for 40 years.

We hasten to add that by no means do all management decision problems lend themselves to treatment in this fashion. The factors to be considered in many important management problem settings simply cannot be usefully quantified. This is especially true at higher, policy levels of management.

On the other hand, most commerce and industry managers serve at lower operational levels. And at the operations level many problem settings involve factors which readily lend themselves to quantification, such as the scheduling of work, budgeting, and price determination.

Managerial Robotics

In general, systems that support managers in dealing with judgment problems amenable to quantification are referred to as decision support systems. The knowledge bases for such systems can frequently be made substantially more useful by invoking the formal analytic methods that assure optimality. Looking further into the future, however, we foresee systems that go beyond merely supporting managers; we predict that managerial robots may become prevalent in the administrative hierarchy of both public and private enterprises.

A managerial robot is the managerial counterpart of an industrial robot. An industrial robot does not simply support a craft worker, but rather replaces that worker in performing that range of functions for which it was designed. In a similar way, a managerial robot does not just support a manager in the performance of some management function, but rather replaces the manager in the performance of that management function. Managerial robots will only become prevalent if they can perform in a fashion significantly superior to the managers they replace, superior both in the quality of the judgments made and the speed with which they are made.

The usefulness of managerial robots will depend on the ability of their designers to accomodate all the features in a judgment setting which real world managers face. A system that simply supports a manager in a particular judgment setting can fall far short of displaying the full range of the manager's skills in dealing with that setting. But a system that replaces the manager can tolerate no such inadequacies.

An Example: The Combinatoric Explosion

For management problems, as well as all other problems, the number of possible decisions, or problem solutions, increases geometrically as the number of variables increases. This combinatoric explosion is a problem for managers, knowledge engineers, and operations researchers. In an effort to contain the many possible combinations of variable values and produce optimal solutions, operations researchers rely on formal algorithmic methods such as the simplex procedure of linear programming and its many variants.

In some operations level management problems, however, the combinatoric explosion proves totally unmanageable when formal analytic tools are employed. Yet recent experience has shown that cleverly designed heuristics can contain the explosion--heuristics which result from in-depth analysis of a type expert managers are not likely to invoke, either because they are ill equipped to perform the required analysis, or because the gains which might be enjoyed from superior judgments are not believed to justify the time and effort which would be required.

In 1984 the coveted Lanchester Prize was awarded to a group producing optimal solutions for a series of scheduling problems. Each used a large mainframe to investigate some 3,000 variables (Crowder, et. al., 1984). In fact, however, many real world scheduling problems, when formulated in such a way that demonstrably optimal solutions can be produced, involve up to four million variables. To produce their optimal solutions the winners of the Lanchester Prize were obliged to ignore features of the problem which are important to managers.

Recently, another group of researchers has attacked large scheduling problems successfully, producing solutions which were

within 98 percent of optimal in minutes on a small micro-computer (Glover, et. al., 1985). By being willing to accept solutions which only approached optimality, this group was able to satisfy all those features in the problem setting which were important to management. Their approach relied entirely on a skillfully designed hierarchy of heuristic rules and yielded a knowledge based system which produces schedules not only faster than an experienced human expert, but also produces schedules which are superior to those constructed by a human expert.

In the latter case, the researchers who addressed the scheduling problem gained insights into the nature of the problem which equipped them to construct a scheduling system which did not simply support the manager in planning the employees' work, but which rather replaced the manager altogether--a true managerial robot.

To summarize, the methods of the operations researchers directly address some of the limitations of human cognitive ability identified by psychologists. Unlike the human expert, an operations research model can sometimes cope with complexity and uncertainty in a systematic, controlled, and explicit manner. We suggest, therefore, that knowledge engineers addressing management problems will want to exploit many of the insights which researchers in this field have gained.

CONCLUSION

We have illustrated some of the approaches knowledge engineers could usefully borrow from psychology and operations research. Research suggests that, for some problems, expert systems developed using these approaches can outperform human experts.

The appropriate set of methods for any knowledge engineering application depends upon the nature of the problem and the experts involved. At present, there is no comprehensive theory to guide the selection of a set of methods for a particular problem/expert combination. Further research is needed to discover how best to integrate new methods into knowledge engineering and how to choose the appropriate set for a particular application. Such research would involve the application of multiple methods to multiple problems, and the

evaluation of results in terms of system development time and system performance criteria such as data demands, speed, and accuracy. Such a research program could contribute to expert system development by allowing knowledge engineers to make more effective use of previous work on modeling judgment and decision processes.

NOTES

1. The distinction between descriptive and prescriptive models is, of course, not as clear as we have drawn it here. It can be argued that the distinction is meaningless because what is descriptive from one perspective may be prescriptive from another. For our purposes, descriptive and prescriptive models are distinguished by the criteria for their validation. Descriptive models are validated according to their ability to reproduce human behavior. Prescriptive models are validated according to the quality of their performance on a specified task.
2. One computer program for judgment analysis is being marketed as an aid to the development of expert systems. Hoffman (1986) describes some advantages of this approach.

REFERENCES

- Allen, G., V. Ciesielski, and W. Bolam (1986). Evaluation of an expert system to forecast rain in Melbourne. Paper presented at the First Australian Artificial Intelligence Congress, Melbourne.
- Armstrong, J. S. (1985). Long-Range Forecasting: From Crystal Ball to Computer, Second Edition. New York: Wiley.
- Camerer, C.F. (1981). General conditions for the success of bootstrapping models. Organizational Behavior and Human Performance, 27, 411-422.
- Crowder, H., E. Johnson, and M. Padberg (1983). Solving large scale 0-1 linear programming problems, Operations Research, 31, 803-834.
- Dawes, R.M., and B. Corrigan (1974). Linear models in decision making. Psychological Bulletin, 81, 95-106.

- Ebbesen, E.B., and V.J. Konecni (1980). On the external validity of decision-making research: What do we know about decisions in the real world? In S. Wallsten (Ed.), Cognitive processes in choice and decision behavior. Hillsdale, New Jersey: Erlbaum.
- Einhorn, H.J., D.N. Kleinmuntz, and B. Kleinmuntz (1979). Linear regression and processing-tracing models of judgment. Psychological Review, 86, 465-485.
- Ericsson, K.A., and H.A. Simon (1984). Protocol Analysis: Verbal Reports as Data. Cambridge, Massachusetts: The MIT Press.
- Fisch H-U., K.R. Hammond, C.R.B. Joyce, and M. O'Reilly (1981). An experimental study of the clinical judgment of general physicians in evaluating and prescribing for depression. British Journal of Psychiatry, 138, 100-109.
- Glover, F., R. Glover, C. McMillan, and S. McMillan (1985). The general employee scheduling problem: An effective large scale solution approach. Proceedings of the Annual Meeting of the Institute for Decision Sciences, 437-441.
- Goldberg, L.R. (1969). The search for configural relationships in personality assessment: The diagnosis of psychosis vs. neurosis from the MMPI. Multivariate Behavioral Research, 4, 523-536.
- Goldberg, L.R. (1970). Man versus model of man: A rationale, plus some evidence, for a method of improving on clinical inferences. Psychological Bulletin, 73, 422-432.
- Hammond, K.R. (1955) Probabilistic functioning and the clinical method. Psychological Review, 62, 255-262.
- Hammond, K.R., T.R. Stewart, B. Brehmer, and D.O. Steinman (1975). Social judgment theory. In M.F. Kaplan and S. Schwartz (Eds.), Human Judgment and Decision Processes. New York: Academic Press.
- Hoffman, P.J. (1960). The paramorphic representation of clinical judgment. Psychological Bulletin, 57, 116-131.
- Hoffman, P.J. (1986). Quasi-AI in the development of expert systems. Los Altos, California, Magic7 Software Co.
- Hogarth, R.M. (1980). Judgment and Choice: The Psychology of Decision. Chichester, England: Wiley.
- Hogarth, R.M. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. Psychological Bulletin, 90, 197-217.
- Kahneman, D., P. Slovic, and A. Tversky, Eds. (1982). Judgment under Uncertainty: Heuristics and Biases. Cambridge, England: Cambridge University Press.

- Kirwan, J.R., D.M. Chaput De Saintonge, C.R.B. Joyce, and H.L.F. Currey. (1983) Clinical judgment in rheumatoid arthritis. I. Rheumatologists' opinions and the development of paper patients. Annals of the Rheumatic Diseases, 42, 644-647.
- Kleinmuntz, B., Ed. (1968). Formal Representation of Human Judgment. New York: Wiley.
- Kruglanski, A.W., N. Friedland, and E. Farkash (1984). Lay persons' sensitivity to statistical information: The case of high perceived applicability. Journal Personality and Social Psychology, 46, 503-518.
- Meehl, P.E. and W.G. Dahlstrom (1960). Objective configural rules for discriminating psychotic from neurotic MMPI profiles. Journal of Consulting Psychology, 24, 375-387.
- Nisbett, R.E., D.H. Krantz, C. Jepson, and Z. Kunda (1983). The use of statistical heuristics in everyday inductive reasoning. Psychological Review, 90, 339-363.
- O'Brien, W.R. (1985). Developing 'Expert Systems': Contributions from decision support systems and judgment analysis techniques. R & D Management, 15, 293-303.
- Oskamp, S. (1967). Clinical judgement from the MMPI: Simple or complex? Journal of Clinical Psychology, 23, 411-415.
- Slovic, P. and S. Lichtenstein (1973). Comparison of Bayesian and regression approaches to the study of information processing in judgment. In L. Rappoport and D.A. Summers, Eds., Human Judgment and Social Interaction, New York: Holt, Rinehart & Winston.
- Stewart, T.R. and M. Glantz (1985). Expert judgment and climate forecasting: A methodological critique of "Climate change to the year 2000". Climatic Change, 7, 159-183.
- Tversky, A., and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. Science, 185, 1124-1131.
- Tversky, A., and D. Kahneman (1981). The framing of decisions and the rationality of choice. Science, 211, 453-458.
- Wallace, H.A. (1923). What is in the corn judge's mind? Journal of the American Society of Agronomy, 15, 300-304.
- Wilsted, W.D., T.E. Hendrick and T.R. Stewart (1975). Judgment policy capturing for bank loan decisions: An approach to developing objective functions for goal programming models. The Journal of Management Studies, 12, 210-215.