SYSTEMS & CONTROL ENCYCLOPEDIA

Theory, Technology, Applications

Editor-in-Chief
Madan G Singh
UMIST, MANCHESTER, UK

Volume 5
M-O

SYSTEMS & CONTROL ENCYCLOPEDIA

Theory, Technology, Applications

VOLUME 5
M–O

Editor-in-Chief

Madan G Singh

University of Manchester
Institute of Science and Technology,
Manchester, UK

PERGAMON PRESS

OXFORD · NEW YORK · BEIJING · FRANKFURT
SÃO PAULO · SYDNEY · TOKYO · TORONTO

U.K.	Pergamon Press, Headington Hill Hall, Oxford, OX3 0BW, England
U.S.A.	Pergamon Press Inc., Maxwell House, Fairview Park, Elmsford, New York 10523, U.S.A.
PEOPLE'S REPUBLIC OF CHINA	Pergamon Press, Qianmen Hotel, Beijing, People's Republic of China
FEDERAL REPUBLIC OF GERMANY	Pergamon Press, Hammerweg 6, D-6242 Kronberg, Federal Republic of Germany
BRAZIL	Pergamon Editora, Rua Eça de Queiros, 346, CEP 04011, São Paulo, Brazil
AUSTRALIA	Pergamon Press Australia, P.O. Box 544, Potts Point, N.S.W. 2011, Australia
JAPAN	Pergamon Press, 8th Floor, Matsuoka Central Building, 1-7-1 Nishishinjuku, Shinjuku-ku, Tokyo 160, Japan
CANADA	Pergamon Press Canada, Suite 104, 150 Consumers Road, Willowdale, Ontario M2J 1P9, Canada

First edition 1987

Library of Congress Cataloging in Publication Data
Main entry under title:
Systems and control encyclopedia.
Includes bibliographies.
1. System analysis—Dictionaries.
2. Control theory—Dictionaries.
3. Systems engineering—Dictionaries.
I. Singh, Madan G.
QA402.S968 1987 003'.03'21 86-15085

British Library Cataloguing in Publication Data
Systems and control encyclopedia: theory, technology, applications.
1. Control theory—Dictionaries I. Singh, Madan G.
003'.03'21 QA402

ISBN 0-08-028709-3 (set)

Printed in Great Britain by A. Wheaton & Co. Ltd, Exeter

SYSTEMS & CONTROL ENCYCLOPEDIA

Theory, Technology, Applications

EDITORIAL ADVISORY BOARD

HONORARY EDITORIAL ADVISORY BOARD

SUBJECT EDITORS

Applications of Control to Aerospace and Aeronautics

M Pélegrin
Centre d'Etudes et de Recherches de Toulouse
Toulouse, France

Applied Robotics, Control Applications to Manufacturing Industries, Flexible Manufacturing and Ergonomics

P Drazan
University of Wales
Institute of Science and Technology
Cardiff, UK

Building Systems

D J Fisk
Department of the Environment
London, UK

Classical Control, CAD of Control Systems, Multivariable Control

N Munro
University of Manchester
Institute of Science and Technology
Manchester, UK

Database Management, Theory and Applications of Distributed Computing for Control

E Gelenbe
Université Paris-Sud
Paris, France

Distributed Parameter Systems, Discrete Systems

J L Lions
Centre National d'Etudes Spatiales
Paris, France

Environmental Measurements, Measurements in Chemistry

Y Sawaragi
Japan Institute of Systems Research
Kyoto, Japan

Environmental Systems and Agricultural Systems

E Halfon
National Water Research Institute
Burlington, Ontario, Canada

Hierarchical Control, Decentralized Control, Model Simplification for Complex Systems Team Theory, Robustness and Sensitivity

A Titli
Centre National de la Recherche Scientifique
Toulouse, France

History of Systems and Control

S Bennett
University of Sheffield
Sheffield, UK

Management Systems, Educational Systems, Health Care Systems, Social Effects of Automation

A P Sage
George Mason University
Fairfax, Virginia, USA

Measurements in the Biological Area

P A Payne
University of Manchester
Institute of Science and Technology
Manchester, UK

Measurements in the Control Area including Intelligent Instrumentation

M G Mylroi
University of Bradford
Bradford, UK

Modelling and Control of Biological Systems, Biomedical Engineering

D A Linkens
University of Sheffield
Sheffield, UK

Pattern Recognition, Fuzzy Systems

M M Gupta
University of Saskatchewan
Saskatoon, Saskatchewan, Canada

Physical Systems Modelling, Identification, Estimation Theory and Applications, Signal Processing

T Kailath
Stanford University
Stanford, California, USA

L Ljung
Linköping University
Linköping, Sweden

Process Control Applications

A J Niemi

R Ylinen
Technical Research Centre of Finland
Espoo, Finland

Realization Theory for Linear and Nonlinear Systems

C Lobry
Nice, France

M Fliess
Centre National de la Recherche Scientifique
Gif-sur-Yvette, France

Y Rouchaleau
Ecole des Mines de Paris
Valbonne, France

Self-Tuning Regulators and Adaptive Control, Stochastic Systems

D W Clarke
University of Oxford
Oxford, UK

Simulation Techniques, Model Simplification

B P Zeigler
University of Arizona
Tucson, Arizona, USA

Stability Theory, General Systems Theory

G J Klir
State University of New York at Binghamton
Binghamton, New York, USA

Static Optimization, Dynamic Optimization, Optimal Control, Game Theory

A Bensoussan
Institut National de Recherche en Informatique et en Automatique
Le Chesnay, France

Technology of Final Control Elements

J P Elloy
Ecole Nationale Supérieure de Mécanique
Nantes, France

Transportation Systems, Communication Systems

A J Fossard
Ecole Nationale Supérieure de l'Aeronautique et de l'Espace
Toulouse, France

Utility Systems and Energy Systems

A Brameller
University of Manchester
Institute of Science and Technology
Manchester, UK

CONTENTS

M

Management: Game Theory

Game theory is the branch of applied mathematics concerned with the modelling of competition by rational players. The seminal work in the field is considered to be *Theory of Games and Economic Behavior* by von Neumann and Morgenstern (1944), which lays out the motivation and structure of the theory. Game theory is a modern discipline, in which new theory and applications are developing. The reader interested in following its progress is referred to the *International Journal of Game Theory*, *Econometrica* and the *Journal of Economic Theory*. Game theory is a mathematical discipline; empirical observation and inference about conflict are generally not admitted. It has had an influence on applied and behavioral disciplines (economics, social psychology, political science and anthropology), which have adopted its modelling approach and the normative conclusions available from simple models (Shubik 1982). This article describes the game-theory approach to modelling and indicates applications along with some results.

A game is a mathematical model of conflict, which may have different representations (Luce and Raiffa 1957, Shubik 1982). There are three representations (forms) that have risen to prominence in the literature: extensive, normal (or strategic) and the characteristic-function forms. The extensive form is the most detailed. It consists of the following:

(a) players (including one called "chance" or "nature" if there are chance or random moves in the game,

(b) a specification of rules of play of the game (order of moves, choices available at each move),

(c) a specification of the information available to each player at each move for that player,

(d) the payoff function (a specification of the value of the game to each player for each set of strategy choices for all the players).

The payoff function is defined using utility functions for the players over outcomes of the game. Such utility functions are shown to exist under axioms suggested by von Neumann and Morgenstern, who demonstrated the existence of utility functions that are unique up to positive linear transformations and showed that the utility of an uncertain situation is just the expected value of the utilities of the events making up the situation. Under these axioms interpersonal comparison of utility by players (e.g., "this outcome is more valuable to me than to you") has no meaning. Many game-theoretic solutions are independent of interpersonal comparisons; however, in many models of cooperative game theory (see below) players' utility functions are assumed linear in money and the effect of interpersonal comparisons are considered.

The extensive form is appropriate for study of the sequential nature of play and is often represented by the "game tree" generated by the sequential moves. When every player knows all of the above elements (a–d), we say that it is a game with complete information. A game with complete information in which, at every move, all previous moves by all players and by nature are known is called a game with perfect information. In this case, the player will know the exact sequence of moves that got him to his current position. In a game with imperfect information, the player will know that some path in some subset of paths led him to his current position, but will not know exactly which one since he has imperfect knowledge of previous moves.

From the extensive form, one can construct the "strategy set" available to each player. A single strategy for a player will be a detailed and complete specification of a choice for each move that the player might be called upon to make for the entire game, conditioned on information available at each move. The strategy set of a player will be the collection of all such strategies. The selection of a single strategy for each player will induce the payoff associated with the play of the extensive-form game under the selected strategies.

A game in normal form consists of the strategy sets for the players and the payoff function, written as a function of a vector of strategy choices, one for each player. The normal-form representation is important because it reduces any game to a game of a single stage with the simultaneous choice of a strategy by all players; the sequential or dynamic nature of play is factored out in this representation. It is the most commonly used form because of its representational economy; it is simply a map from strategies to payoffs. Often, extensive- or normal-form games are described by the number of players, as, for example, an n-person normal-form game. A class of interest in normal- and extensive-form games are zero-sum games: for all strategy choices for the players the sum of payoffs to all players is a constant. Often, players are allowed to choose strategies which are themselves randomizations of strategies, which are referred to as mixed strategies. If a game allows players to coordinate their strategies, we refer to the coordinated strategy as a correlated strategy, and if it admits mixed strategies, as a mixed correlated strategy.

The characteristic-function form assigns to each subset of players, called coalitions, a set of payoffs to the members of that coalition. The characteristic-function form is used to study behavior in situations where cooperation between players is possible. If the utility of payoffs to the players are transferable and linear in money, the value of a coalition is a number corresponding to the maximum total utility of that coalition, defined in one of two ways: the "α" or "β" form. The value under the α form is the total utility that the coalition can guarantee its members by using a correlated strategy of members, no matter what strategy the players outside the coalition adopt. The β form is the total utility that the coalition cannot be prevented from getting, no matter what the players outside the coalition do. A similar definition of the forms can be made in the case of nontransferable utility.

How the players behave and what a reasonable outcome of the game is are specified by the solution concept. Solution concepts are divided into two types: cooperative and noncooperative. In noncooperative games, binding agreements between players are not possible and players recognize that any agreement to coordinate play and share value cannot be enforced. Players can agree to coordinate their strategies, but the agreement will only stand if each player cannot strictly improve its payoff by "double-crossing" its partners: any agreement must be self-serving. The basis of noncooperative theory for both the extensive and the normal form is the Nash equilibrium, a combination of strategy choices such that each player's strategy is the best response to the other players. Since the set of Nash equilibria for a game may be composed of more than one element, a host of refinements have arisen: subgame perfectness and sequential, perfect, proper and persistent equilibria. Each is based on intuitive ideas of what properties a noncooperative equilibrium should have.

Cooperative solution concepts allow binding agreements between players. Solution concepts, defined for the characteristic-function form can be divided into two categories: equitable solutions and unobjectionable solutions. For equitable solutions, axioms describing properties of "fair" solutions characterize the Nash bargaining solution, the Kalai–Smorodinsky solution and the Shapley value. The nucleolus solution for games of characteristic-function form with linearly transferable utility is an equitable solution which assigns a payoff to each player such that the minimum difference between the imputed value and the characteristic-function value for all coalitions is maximized. Unobjectionable solutions, such as the core, the von Neumann–Morgenstern solution set and the bargaining set, all describe feasible payoff vectors to the players that no coalition can "block" or find objection to in some sense.

An important, relatively recent development in game theory is Harsanyi's approach to games with incomplete information (Myerson 1984). Suppose that for an extensive-form game the payoff functions of the players are dependent upon some state variable whose value is unknown. Each player independently and privately observes signals that give information about the state through conditional probability statements on states given signals. Harsanyi assumes that the prior distribution on states and the conditional joint distribution on players' signals given states are known by all. The game becomes one of complete but imperfect information. Harsanyi defines the noncooperative solution concept of a Bayesian equilibrium for a game with incomplete information as the vector of strategies for each player conditioned upon individual signals, such that each player's strategy choice (given his individual signal) maximizes his expected utility (given the conditional distribution of the others' private signals and therefore strategy choices).

1. Applications from Noncooperative Theory

Game theory has made a strong contribution to areas of managerial interest in the field of economics, particularly in the areas of oligopolistic behavior and the economics of decentralized organizations (Schotter and Schwodiauer 1980).

1.1 Oligopolistic Competition

The economic theory of oligopolistic behavior studies competition of the few and applies to all markets where perfect competition cannot be assumed. Game theory's contribution to the existing literature has been the study of the use of price, quantity, location, advertising, product spectrum, product innovation, production capacity and reputation as strategic variables in a variety of models to explain competitive behavior. Recent work has focused on issues of incomplete information in such models and the solution of the models in repeated situations, since rarely do firms play against each other just once. The game-theory approach has yielded exact solutions to abstract market models and has amassed a body of positive results about the use of the above variables in real markets.

1.2 Decentralization

Intrafirm decision making focusing on issues of coordination and incentives has received considerable scrutiny by game theorists and economists (Jennergren 1980). There are a number of threads in this research, most of which model the decentralized decision process as that of an n-person normal-form game and seek properties of decisions corresponding to the Nash equilibrium. Marschak and Radner (1972) in their theory of teams studied decision making when all agents hold the same objective function but differing information about the state of the system including others' decisions, and examine optimal information structures and decision rules for such organizations. Groves and Loeb studied the problem of group decision making, closely related to the economic problem of allocating public goods, with decision makers with differing individual utility functions on some space of group alternatives one of which must be selected (Green and Laffont 1979). Their analysis addresses directly the problem of allocation of resources and other common decisions made by headquarters of multidivisional firms. Each agent is assumed to have information about his own utility function, but not about the others'. The agents will communicate to a center (headquarters) that interprets the agents' messages and decides on a group decision through some mechanism. The class of incentive schemes (functions specifying rewards to the players) that yields messages that contain truthful revelation of utilities as a dominant strategy (truth telling is the best response to any strategy of opponents) of the induced game is characterized and has been found to be unsatisfactory. In these schemes, some agents can be individually worse off with the decision than they were before. No mechanism can tax the agents to pay for the selected alternative such that the tax collected exactly equals the cost. The total profit of all agents (the firm's total profit) may not be maximized, and the mechanism is manipulatable by coalitions, that is, coalitions may not find truth telling to be a dominant strategy.

A related literature concerns the allocation of resources in hierarchical organizations, where the firm's total profit is the sum of divisional profit, the division manager's goal is to maximize his individual profit, and divisions are interrelated only by resource constraints. Under uncertainty about divisional profit functions, the headquarters of such a firm wishes to decide how to allocate a fixed quantity of resources to maximize total profit. This is done through an interchange of information about the marginal utility of resources to divisions, who myopically respond to transfer prices on resources with messages of demand quantities that maximize divisional profit. Adjustment of transfer prices continue as a function of quantity messages, until total demand equals supply. Under very restrictive conditions, convergence will occur. Although this approach is not explicitly game-theoretic, we may view the adjustment process as a multi-stage game, where the payoff to the division is the division's profit when the process stops. Truthful revelation of demands are a dominant strategy equilibrium of the game and yield, in the final stage, the profit-maximizing allocation for the entire firm.

Recent interest has been in decentralized decision making under incomplete information when headquarters' profit depends on information from divisions and on decisions made locally by divisions. This is known as the principal-agent problem. Researchers have studied incentive schemes that maximize headquarter's profit such that truthful relevation of information by divisions and optimal divisional decisions are a Bayesian equilibrium of the induced game (Myerson 1984). Such schemes ensure that the firm is maximizing profit to the fullest possible extent, even though the truthfulness of information sent cannot be verified (the problem of "adverse selection"), nor can the local decisions be monitored by headquarters (the problem of "moral hazard"). In these models, the players' differing attitudes to risk are explicitly considered. Such models have been used to analyze optimal employment contracts and insurance policies where information from applicants for a job (or insurance contract) cannot be checked and the job holder's performance (or insurance holder's behavior with the insured property) cannot be perfectly monitored or the checking/monitoring is prohibitively expensive.

1.3 Auctions and Bidding

A considerable literature has developed on auction mechanisms, especially the English ascending oral auction (the usual oral auction), the first-price sealed bid auction (the high bidder wins and pays its bid), the Dutch auction (the auctioneer announces a sequence of descending prices, the first bidder who responds wins and pays the announced price) and the second-price sealed bid auction (the high bidder wins and pays the bid of the second highest bidder). Optimal bidding in such markets has also been studied. A comprehensive survey of current work and a bibliography are available in Engelbrecht-Wiggans (1980). Two modelling approaches predominate in the literature: decision analysis, where a bidder is considered the only strategic player, and the game-theory approach, where all bidders are strategic. In studying bidding behavior, the first approach stresses forecasting the distribution of bids facing the bidder, with an optimal bid set on the basis of expected utility, the second approach studies properties of the Bayesian equilibrium of the induced game with incomplete information. The game-theory approach has generated a raft of qualitative and quantitative conclusions about optimal bidding and the revenue generating properties of auctions (Milgrom and Weber 1982).

An important distinction has developed about the bidders' valuation of auctioned objects and its effect on models: private values as distinguished from common values. In the first case, each bidder's estimate of value is private and others' valuations do not affect his. A bidder is assumed to have an estimate of his own valuation and has a probability distribution on the others' valuations. This model would be appropriate for auctioned ojbects where others' feelings about the quality of an object are unimportant, there is no quantity uncertainty, and the object will be consumed or not brought back to the market for resale for a long period. In the common values model, a bidder's estimate of value is affected by the valuation of others. This model is appropriate for bidding situations in which the object auctioned will be resold, or where there is quantity or quality uncertainty. An auction involving a house to be bought for quick resale, or mineral recovery rights, or a rare painting would be appropriate for the common values model. The distinction between models is important because, in the common values case, a firm's strategy in bidding imparts information to a competitor not only about the bidders' valuation but also the value of the object to the competitor. The information content of bids cannot be ignored. For example, typically in bidding for oil reserves, there is considerable uncertainty in the quantity of oil recoverable. A firm that bids and wins probably will have overestimated the size of the recoverable reserve unless he has taken into account what winning the bid says about the others' estimates.

It has been demonstrated that for the private values model, the English, first-price sealed bid, second-price sealed bid and Dutch auctions will all deliver the same expected revenue to the seller, in the absence of risk aversion of the bidders. With risk aversion, the first-price sealed bid and Dutch auctions will yield higher expected revenues than the others. For the common values model with risk neutrality of bidders; the English auction generates higher expected revenues than the second-price sealed bid, which is in turn superior to the first-price sealed bid and the Dutch auction. The problem of optimal auction design has yielded results on alternate auction forms, not of the above types, that maximize seller revenue in private values models. Research has also been done on the effect that entry fees and reserve prices have on expected revenues.

1.4 Games against Nature

Games involving one strategic player in situations with uncertainty have been studied in statistical decision theory for situations where the opponent is not a strategic player (for example, the stock market facing the strategic investor) and where the opponent is an active player but his motivations and actions are not modelled. The two routes in analysis are

(a) decision theory: assume a probability distribution for opponents actions, and act optimally against that distribution;

(b) game theory: assume the conservative stance that the opponent's interests are strictly opposed to yours and model the players' decision problem as a two-person zero-sum game, a model form well studied and understood.

Applications of the game-theoretic approach have included capital budgeting, sampling and investment in securities options. Solution criteria for these types of games have taken the form of either a "minimax" solution (the player plays the strategy that maximizes his minimum winnings), which corresponds to the player's Nash equilibrium strategy of the game; Savage's "minimax regret criterion" (the minimax strategy for the new zero-sum game generated by defining the payoff for each strategy pair to be the difference between the player's payoff and the best payoff possible for the opponent's strategy choice for the original game); or Hurwicz's "pessimism–optimism index criterion" (where the player plays his minimax strategy against a weighted combination of the best and worst possible outcomes). See Luce and Raiffa (1957) for a discussion.

2. Applications from Cooperative Theory

2.1 Bargaining, Arbitration and Mediation

Game theory has attempted to cast light on the process of bargaining between individuals by viewing the process as a cooperative one where the players involved must agree on an assignment of utility among themselves. Outcomes of the bargaining process are in terms of utility assignments, not in terms of strategy choices. The criteria of this assignment have been efficiency (there is no other assignment of utility to the players through correlated strategies where all do at least as well and at least one player is doing strictly better) and equity (the fairness of the assignments, considering the initial utility assignments of the players and their power, as measured by their strategic possibilities). Bargaining theory is an axiomatic theory; a set of axioms are assumed and the properties of the solutions obeying the axioms are characterized (Roth 1979). One bargaining solution is distinguished from another by the axioms of fairness underlying it. An alternative interpretation of this process has been that a third party, an arbitrator, must decide upon an assignment, based upon axioms of efficiency and equity, which the players are bound to obey. (Mediation is viewed as a third party process to aid the bargainers, which is not binding.) John Nash proposed one set of axioms that an equitable bargaining procedure should have: symmetry, independence of irrelevant alternatives, individual rationality, efficiency and the requirement that the solution be independent of positive linear transformations on the player's utility functions. With these axioms he demonstrated the existence of a unique bargaining solution. Given an n-person normal-form game, the Nash bargaining solution assigns a unique solution: the vector $(v_1, v_2, \ldots, v_n)$ of utility to the players that maximizes the product $(v_1 - u_1)(v_2 - u_2) \ldots (v_n - u_n)$ over the set of utility n-tuples generated by the set of mixed correlated strategies of the players, where $(u_1, u_2, \ldots, u_n)$ are minimax levels for the players. Other axiomatic bargaining solutions are based on the work of Raiffa and of Kalai and Smorodinsky. What happens in a bargaining situation when the players fail to reach an agreement and the effect that this has on the bargaining process have been modelled by Nash. Unless the players can agree on a solution, they will resort to a disagreement or "threat" strategy. Nash's variable threat solution assigns to each player the utility of the n-person Nash equilibrium of the game generated by the Nash bargaining solution with the $(u_1, u_2, \ldots, u_n)$ vector replaced by the payoff vector of the threats. Thus, the bargainers seek an equilibrium in threats.

Nash's solution assumes complete information about the players' strategies and payoffs. Harsanyi and Selten (1972) and Myerson (1979) solve the problem of bargaining under incomplete information by defining the solution to be utility assignments that are uniquely chosen from the convex hull of payoffs of the Bayesian equilibria of a well-defined bargaining game. Another approach is that of bargaining when payoff functions of others are unknown (Sobel 1981). If payoff functions are individually reported to an arbitrator who applies the Nash bargaining solution or the Kalai–Smorodinsky solution, the players will play a game with payoff function reports as strategies. For the case of bargaining over a single commodity, the Nash equilibrium of the reporting game will have all players report linear, risk-neutral payoff functions and this will be a dominant strategy. The implication of this result is, when bargaining, to represent your utility function as being risk-neutral.

To date, the game-theory approach has not well integrated the factors of history, psychology and institutional structure into the bargaining problem. Thus, game-theoretic results can only serve as a guide, a "standard," for any specific bargaining situation. See Walton and McKersie (1965) for a discussion of how game-theoretic and extra-game-theoretic factors come into play in a collective bargaining process.

2.2 *Allocating Costs*

Cooperative theory has been applied to the description of various methods of cost allocation for intrafirm sharing of the cost of a commonly used facility and for societal sharing of the cost of government projects. The solution concepts of the core, the Shapley value and the nucleolus have been used for the allocation of intrafirm costs and costs of air pollution regulation, internal telephone billing rates, airport use, public utility pricing and water resources development (Shubik 1982).

These solution concepts are defined for characteristic-function games with linearly transferable utility, and are defined, in the current case, for agents whose utilities are just total profit. The characteristic-function value of any coalition is its net increase in profit using the facility and paying for it itself (assuming the net increases is positive, otherwise the value is zero). The core is the set of allocations for the players (each summing to the value of the coalition of all players) such that each coalition's allocation is at least as great as its value. Thus, each coalition's net return is at least as great as it can guarantee itself. A problem with the core is that it may not exist, and if it does, it generally will not be unique.

The Shapley value always exists and is unique. It is defined as the net profit allocation such that each agent's allocation is exactly its expected marginal contribution to a randomly selected coalition. The allocation given by the Shapley value sums to the value of the coalition of all players. Like the core, the Shapley value guarantees that each player in the game enjoys at least as much profit as if he were standing alone. A problem of the Shapley value, however, is that the allocation generated by it may not have the above property for a coalition.

An allocation scheme which always exists, is unique and attempts to minimize coalition complaints to allocations is the nucleolus. It has been pointed out that this scheme is like a Rawlsian cost-share scheme, where the payoff to the worst-off agent is maximized. A problem is that an individual agent or coalition may receive an allocation less than it can guarantee itself.

2.3 *Value of a Block of Shares*

Another application of the Shapley value is to measure the power of players in a voting game where a fixed percentage of votes rules. This approach has been generalized to apply to the value of a block of shares in corporate takeovers (Shubik 1982).

See also: Games with Infinitely Many Players; Games, Noncooperative; Games, Repeated; Differential Games: Introduction; Nash Equilibrium; Decision Analysis; Human Judgment and Decision Rules

Bibliography

Engelbrecht-Wiggans R 1980 Auctions and bidding models: A survey. *Manage. Sci.* **26**, 119–42

Green J R, Laffont J-J 1979 *Incentives in Public Decision Making*. North-Holland, Amsterdam

Harsanyi J, Selten R 1972 A generalized Nash solution for two-person bargaining games with incomplete information. *Manage. Sci.* **18**, 80–106

Jennergren L P 1980 On the design of incentives in business firms—A survey of some research. *Manage. Sci.* **26**, 180–201

Luce R D, Raiffa H 1957 *Games and Decisions*. Wiley, New York

Marschak J, Radner R 1972 *Economic Theory of Teams*. Yale University Press, New Haven, Connecticut

Milgrom P R, Weber R J 1982 A theory of auctions and competition and the bargaining problem. *Econometrica* **47**, 61–73

Myerson R B 1979 Incentive compatibility and the bargaining problem. *Econometrica* **47**, 61–73

Myerson R B 1984 Bayesian equilibrium and incentive—Compatibility: An introduction. In: Hurwicz L, Schmeidler D, Sonnenschein H (eds.) *Social Goals and Social Organization, Essays in Memory of Elisha A Pazner*. Cambridge University Press, Cambridge

Roth A E 1979 *Axiomatic Models of Bargaining*. Springer, Berlin

Schotter A, Schwodiauer G 1980 Economics and the theory of games: A survey. *J. Econ. Lit.* **18**, 479–527

Shubik M 1982 *Game Theory in the Social Sciences*. MIT Press, Cambridge, Massachusetts

Sobel J 1981 Distortion of utilities and the bargaining problem. *Econometrica* **49**, 597–619

von Neumann J, Morgenstern O 1944 *Theory of Games and Economic Behavior*. Wiley, New York

Walton R E, McKersie R B 1965 *A Behavioral Theory of Labor Negotiations*. McGraw–Hill, New York

P. J. Lederer

Management: Hypothesis Testing

The two major areas of statistical inference are the testing of hypotheses and the estimation of system states and parameters. The former area, hypothesis testing, is the subject of this article. We will develop some aspects of classical decision theory, and describe simple applications to communication and control sciences.

There are many situations in which hypothesis testing (often called detection theory in the communication theory literature) is applicable. A radar return is observed, and the presence or absence of a target is to be determined; from a smear of human tissue we attempt to determine whether a patient has cancer; from fluctuations in activity of a particular stock on one of the exchanges, we decide whether or not to buy some shares of the stock.

In each case we use choose an answer yes or no, and we refer to these two choices (yes or no) as hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$. It is possible to have more than two hypotheses and to attempt to decide which one of N hypotheses to accept.

1. Purpose of Hypothesis Testing

In the analysis phase of the systems process, it is often desirable to assess the validity of various assumptions concerning systems operation. In any situation where hypotheses or assumptions are used, there is the possibility of accepting a false hypothesis or rejecting a true hypothesis. These two situations are called Type I and Type II errors. Figure 1 indicates Type I and Type II errors for a hypothetical situation. In the absence of a

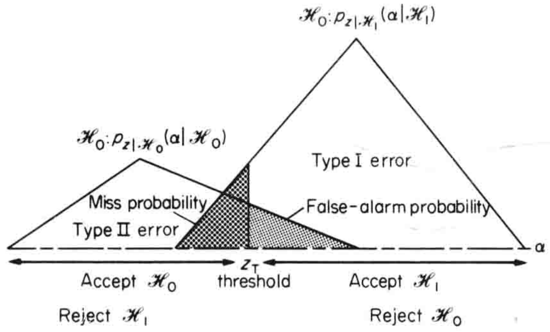

Reality	Decision: Accept $\mathcal{H}_0$	Decision: Reject $\mathcal{H}_0$
$\mathcal{H}_0$ true	Correct decision	Type I error (probability α)
$\mathcal{H}_0$ false	Type II error (probability β)	Correct decision

Figure 1
Probability densities for illustration of Type I and Type II errors. A single sample is made, and if the result is less than z_T, then $\mathcal{H}_0$ is accepted; otherwise $\mathcal{H}_1$ is accepted

hypothesis testing method, the probability of a Type I or Type II error is unknown. Hypothesis testing provides measures that enable determination of whether too much significance is attached to results which occur through pure chance. It is not generally possible to reduce the probability of a Type I or Type II error to zero. However, either but not both of these probabilities can be reduced to the extent that they are virtually eliminated. This process is known as establishing a level of significance.

Hypothesis testing is used when assumptions must be evaluated, and where it is not possible to study the entire "population" of values about which these assumptions have been made. In its simplest form, two hypotheses are defined: the null hypothesis $\mathcal{H}_0$ and the alternate hypothesis $\mathcal{H}_1$. Typically, the alternate hypothesis is developed first. It usually concerns the positive aspect of an issue and may come directly from the issue formulation elements. For example, the alternate hypothesis might be: "There *has been* a reduction in energy use in *London* over the past twelve months." It is important to note that the basis of the hypothesis testing method is the null hypothesis. Type I and Type II errors are based on rejection of, or inability to reject, the null hypothesis. In other words, the alternate hypothesis will only be accepted if the null hypothesis is rejected. If the null hypothesis cannot be rejected then no conclusion is possible concerning the alternate hypothesis. The null hypothesis and the alternate hypothesis must be selected so that data supporting the alternate hypothesis is such that it can potentially cause the null hypothesis to be rejected. $\mathcal{H}_1$ is designed so that it is more likely to be true if $\mathcal{H}_0$ is rejected. The following statement would be inappropriate as a null hypothesis to accompany the previous $\mathcal{H}_1$. "There *has not been* a reduction in energy use in the *UK* over the past twelve months." Results of a sample taken to determine the validity of this hypothesis would not provide meaningful information about $\mathcal{H}_1$. An appropriate null hypothesis is the negative of the alternate hypothesis. "There *has not been* a reduction in energy use in *London* over the past twelve months" is a good null hypothesis to use in conjunction with the $\mathcal{H}_1$ stated above.

Following the selection of the two hypotheses, specification of the level of significance is needed. The value of this level is arbitrary, and is specified in accordance with the decisionmaker's desires. The level of significance α is compared to a statistical analysis on the results of an experiment performed on a random sample of the population. This "population" is the same set of values for all tests. Statistical analysis of the test results yields the probability of finding the results of that test if $\mathcal{H}_0$ were actually true. If the statistical value is smaller than the present level of significance, it is concluded that $\mathcal{H}_0$ should be rejected and $\mathcal{H}_1$ accepted. If the statistical value, or likelihood, is greater than α it is concluded that the sample information does not provide sufficient evidence to reject $\mathcal{H}_0$. The deviation from data that support $\mathcal{H}_0$ found in the sample could be due solely to random factors. It is desired to reduce this possibility. Tests are generally designed to keep the possibility of making Type I errors small. Since the null hypothesis usually states what is believed to be true in the absence of data, it is generally appropriate to structure the problem such that supporting data will tend to cause rejection of the null hypothesis and acceptance of the alternate hypothesis. Often it is more costly to make a Type I error than a Type II error of the same relative amount, and if this is the case we try to avoid making that type of error.

Selection of an appropriate test statistic is based on the population being studied. Normal distributions, chi-square, t or F distributions are typical choices and are picked as a function of the degrees of freedom or continuity of range in the possible outcomes. For example, the possible outcomes of the roll of a pair of dice has fewer degrees of freedom than the outcome of selecting a colored jellybean from a jar of 2000 beans of 20 different colors. Proper selection of the distribution is based on the statistical properties (mean, variance, etc.) of the population. However, if these are not known, sample results may be used to calculate these values. This complicates the computation of error probabilities considerably.

Design of the sample size is often a function of how much freedom the analyst has to choose the sample size. Statistical methods for computing the minimum sample size are available. The sample size is based on the minimum size necessary to guarantee that certain constraints will be met. The parameter α is the probability that a Type I error will be made, and β the probability that a Type II error will be made. Although β cannot be calculated unless $\mathcal{H}_1$ is an exact hypothesis, it is known that, for a specified α, the parameter β will decrease as the sample size increases. This means that the possibility of properly rejecting a false hypothesis increases as more information becomes available.

Extreme care must be taken to ensure that all samples are taken in a truly random fashion. The validity of the statistical outcome depends heavily on the randomness of the sample, and this indicates the considerable importance of randomness in the collection process.

Once the mean, variance and other relevant aspects of the test have been calculated, the value "p" must be calculated. Here p is the likelihood that a particular test result, given that $\mathcal{H}_0$ is true, is found in a random sample. It also represents the fraction of times that the test results achieved in the test can be expected among a very large number of similar tests, given that $\mathcal{H}_0$ is in fact true. The analyst compares the value of p to the preset α. If p is smaller than α, it is concluded that this particular test result is so unlikely if $\mathcal{H}_0$ is true that $\mathcal{H}_0$ must be rejected. If p is greater than α then it is concluded that the deviation in the test result from $\mathcal{H}_0$ could be accounted for by random chance. There is, in this case, not enough evidence to reject $\mathcal{H}_0$.

Hypothesis testing is useful in any situation where it is desirable that hypotheses be accepted or rejected on the basis of statistical information. It is not uncommon that the issue formulation and analysis steps of the systems process do not initially result in complete and precise information about one or more aspects of system operation. It is often useful to have a method to evaluate assumptions about the system, rather than basing an important decision merely on intuition. It is to these ends that hypothesis testing is most useful. Hypothesis testing can be valuable in the system implementation or operation when an action depends on a judgement about the actual state of a system, and where complete information about system operation is not attainable.

2. *Elementary Mathematics for Hypothesis Testing*

Each hypothesis in a given problem results from a source which generates one of the hypotheses as an output. Data to enable acceptance or rejection of the hypothesis is not observed directly, or there would be no decision problem. A probabilistic transition mechanism "separates" the hypothesis from our observations. This device knows which hypothesis is true and generates a point or points in observation space according to some probability law. We have access to the observation space, and form a decision on the basis of a decision rule as to which hypothesis to accept. This decision is based upon a knowledge of the *a priori* probability of the various hypotheses and of the conditional probabilities inherent in the probabilistic transition mechanism. If the (*a priori*) probability that $\mathcal{H}_0$ will occur is $P_{\mathcal{H}_0}$ and the probability of $\mathcal{H}_1$ is $P_{\mathcal{H}_1}$, then $P_{\mathcal{H}_0} + P_{\mathcal{H}_1} = 1$, since one of the hypotheses $\mathcal{H}_0$ or $\mathcal{H}_1$ must be true. We make a single observation z corrupted by noise (which may have different statistical characteristics, depending upon which hypothesis was true), and attempt to make a decision from the hypotheses

$$\mathcal{H}_0: z = v_0$$

$$\mathcal{H}_1: z = 1 + v_1$$

Knowledge of the probabilistic transition mechanism is equivalent to a knowledge of the probability density functions associated with v_0 and v_1. In much of the work in systems and control v_0 and v_1 are called measurement noise. Figure 2 illustrates the elements of the binary hypothesis testing problem.

Two kinds of errors are possible in simple binary decision problems, as we have stated. We may accept

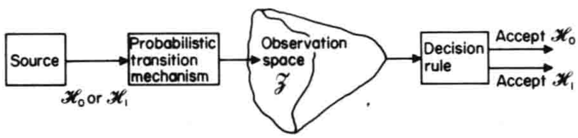

Figure 2
Elements of elementary hypothesis testing problem

$\mathcal{H}_0$ when it is false, or we may accept $\mathcal{H}_1$ when it is false. In the radar problem, where $\mathcal{H}_0$ corresponds to the absence of a target and $\mathcal{H}_1$ corresponds to the presence of one, acceptance of $\mathcal{H}_0$ when it is false is called a "miss," and the probability of doing this is called the miss probability P_M. One minus the miss probability is called the detection probability, $P_D = 1 - P_M$. Accepting $\mathcal{H}_1$, deciding that a target is present when $\mathcal{H}_1$ is false and a target is not present is called a false alarm and the associated probability is the false-alarm probability P_F.

These ideas may be illustrated in a simple graphical fashion. An alternative statement of the problem is that under $\mathcal{H}_0$, we have $p_z(\alpha) = p_{z|\mathcal{H}}(\alpha|\mathcal{H}_0)$, and under $\mathcal{H}_1$ we have $p_z(\alpha) = p_{z|\mathcal{H}}(\alpha|\mathcal{H}_1)$. Thus the object of the hypothesis testing problem is to accept one of the two density functions $p_{z|\mathcal{H}_0}$ or $p_{z|\mathcal{H}_1}$ as being "most representative" of the density of a given population. These densities may appear as illustrated in Fig. 1. Assume a threshold at $\alpha = z_T$, so that if a single observation $z = \alpha_1$ is greater than z_T, we accept $\mathcal{H}_1$ and if it is less than z_T, we accept $\mathcal{H}_0$. In a more realistic situation the parameter α is defined as a sufficient statistic, and may be made up of several observations rather than obtained from a single observation. Here, we find it convenient to consider α as a scalar observation.

We determine false-alarm and miss probabilities for a simple single observation case. From Fig. 1 and the foregoing reasoning the miss probability is

$$P_M = \int_{-\infty}^{z_T} p_{z|\mathcal{H}}(\alpha|\mathcal{H}_1)\,d\alpha \tag{1}$$

while the false-alarm probability becomes

$$P_F = \int_{z_T}^{\infty} p_{z|\mathcal{H}}(\alpha|\mathcal{H}_0)\,d\alpha \tag{2}$$

The miss probability may be made as small as desired at the expense of the false-alarm probability. In a practical situation one of the probabilities, such as P_F, may be fixed, and a test is selected. This might be the threshold z_T selected, so that P_M is minimized. This criterion, known as a Neyman–Pearson criterion, is trivial in this particular situation, since fixing the false-alarm probability fixes the threshold, and hence the miss probability. In a more realistic situation there would be more than one observation, and thus ample opportunity to constrain the false-alarm probability by using Lagrange multipliers without simultaneously fixing the miss probability. Hence, we would obtain a minimum of

$$J_{NP} = P_M + \lambda[P_F - \gamma] \tag{3}$$

where λ is a Lagrange multiplier, and γ is the desired false-alarm probability. We set the derivative of J_{NP} with respect to z_T equal to zero to obtain, by using Eqns. (1) and (2) such that we have a likelihood-ratio test for an observation $z = \alpha$,

$$\frac{p_{z|\mathcal{H}}(\alpha|\mathcal{H}_1)}{p_{z|\mathcal{H}}(\alpha|\mathcal{H}_0)} \underset{\text{accept } \mathcal{H}_0}{\overset{\text{accept } \mathcal{H}_1}{\gtrless}} \lambda \tag{4}$$

which says that we accept $\mathcal{H}_1$ if the ratio is greater than λ and accept $\mathcal{H}_0$ if the ratio is less than λ when evaluated at $z = \alpha$. We adjust the variable λ, the Lagrange multiplier, so that $P_F = \lambda$.

If a false alarm is as serious as a miss, an appropriate criterion might be to minimize the miss probability plus false-alarm probability. Thus we minimize

$$P_{F+M} = P_M + P_F = \int_{-\infty}^{z_T} p_{z|\mathcal{H}}(\alpha|\mathcal{H}_1)\,d\alpha + \int_{z_T}^{\infty} p_{z|\mathcal{H}}(\alpha|\mathcal{H}_0)\,d\alpha \tag{5}$$

by a proper choice of z_T. We take the derivative of the foregoing expression with respect to z_T and set it equal to zero to obtain the threshold as the value for which the two densities are equal. If $z = \alpha$, then $p_{z|\mathcal{H}_1} > p_{z|\mathcal{H}_0}$, and we accept hypothesis $\mathcal{H}_1$. In this case we have a likelihood-ratio test such that, since λ is in effect equal to 1,

$$\frac{p_{z|\mathcal{H}}(\alpha|\mathcal{H}_1)}{p_{z|\mathcal{H}}(\alpha|\mathcal{H}_0)} \underset{\text{accept } \mathcal{H}_0}{\overset{\text{accept } \mathcal{H}_1}{\lesseqgtr}} 1 \tag{6}$$

Hypothesis testing problems are generally much more complex than the simple scalar observation problem presented here. Often it is unrealistic to assume that false alarms and misses are equally serious. There may be different costs associated with these errors, and there may be costs associated with correct decisions. The Bayes test, or Bayes risk criterion, is used to treat problems of this type. We briefly examine this criterion here.

The four courses of action in testing hypotheses against single alternatives and their associated costs are

C_{00}: cost of accepting $\mathcal{H}_0$ when $\mathcal{H}_0$ is true
C_{01}: cost of accepting $\mathcal{H}_0$ when $\mathcal{H}_1$ is true
C_{10}: cost of accepting $\mathcal{H}_1$ when $\mathcal{H}_0$ is true
C_{11}: cost of accepting $\mathcal{H}_1$ when $\mathcal{H}_1$ is true

C_{00} and C_{11} represent costs associated with correct decisions, whereas C_{01} and C_{10} represent costs associated with incorrect decisions.

It is not unrealistic to associate "costs" with correct decisions. In an application of decision theory to investment, for example, the cost associated with not buying a stock, when the correct decision is *not* to buy a stock, is that the money which could be invested in the market for potentially large returns is not being invested, or is being invested at much lower rates of return. In a similar way, there is "cost" associated with purchasing a certain stock when the correct decision should be to buy that stock, since the money invested is now "risk capital,"