

Game Theory

ANDREW M. COLMAN

Volume 2, pp. 688–694

in

Encyclopedia of Statistics in Behavioral Science

ISBN-13: 978-0-470-86080-9

ISBN-10: 0-470-86080-4

Editors

Brian S. Everitt & David C. Howell

© John Wiley & Sons, Ltd, Chichester, 2005

Game Theory

Background

Game theory is a formal theory of interactive decision making, used to model any decision involving two or more decision makers, called *players*, each with two or more ways of acting, called *strategies*, and well-defined preferences among the possible outcomes, represented by numerical *payoffs*.

In the theory, a player can represent an individual human decision maker or a corporate decision-making body, such as a committee or a board. A (pure) strategy is a complete plan of action, specifying in advance what *moves* a player would make – what actions the player would perform – in every contingency that might arise. If each player has only one decision to make and the players decide simultaneously, then the concept of a strategy coincides with that of a move, but in more complicated cases, a strategy is a comprehensive plan specifying moves in advance, taking account of all possible counter-moves of the coplayer(s). A *mixed strategy* is a probability distribution over a player's set of pure strategies. It is usually interpreted as a strategy chosen at random, with a fixed probability assigned to each pure strategy, as when a player tosses a coin to choose between two pure strategies with equal probabilities. In Bayesian interpretations of game theory, initiated by Harsanyi [12], a mixed strategy is sometimes interpreted as a coplayer's uncertainty about a strategy choice. Payoffs represent players' von Neumann–Morgenstern utilities, which are (roughly) their true preferences on an interval scale of measurement, as revealed by the assumedly consistent choices that they make in lotteries in which the outcomes have known probabilities assigned to them. A player's *payoff function* is a mapping that assigns a specific payoff to each outcome of the game.

The conceptual groundwork of game theory was laid by Zermelo, Borel, von Neumann, and others in the 1920s and 1930s (*see* [10]), and the first fully developed version of the theory appeared in *Theory of Games and Economic Behavior* by von Neumann and Morgenstern [31]. The theory began to have a significant impact on the behavioral and social sciences after the publication in 1957 of a more accessible text entitled *Games and Decisions* by Luce and Raiffa [18]. The early game theorists considered

the chief goal of the theory to be that of prescribing what strategies rational players *ought* to choose to maximize their payoffs. In this sense, the theory, in its classical form, is primarily *normative* rather than *positive* or *descriptive*. An additional rationality assumption, that people generally try to do the best for themselves in any given circumstances, makes the theory relevant to the empirical behavioral sciences and justifies experimental games (*see* section titled Experimental Games below); and in evolutionary game theory, the rationality assumption is replaced by replicator dynamics or adaptive learning mechanisms (*see* section titled Evolutionary Game Theory).

Basic Assumptions

In conventional decision theory, rational choice is defined in terms of maximizing expected utility (EU), or subjective expected utility (SEU), where the objective probabilities of outcomes are unknown (*see* utility theory). But this approach is problematic in games because each player has only partial control over the outcomes, and it is generally unclear how a player should choose in order to maximize EU or SEU without knowing how the other player(s) will act. Game theory, therefore, incorporates not only rationality assumptions in the form of expected utility theory, but also common knowledge assumptions, enabling players to anticipate one another's strategies to some extent, at least. The standard common knowledge and rationality (CKR) assumptions are as follows:

CKR1 (common knowledge): The specification of the game, including the players' strategy sets and payoff functions, is *common knowledge* in the game, together with everything that can be deduced logically from it and from the rationality assumption CKR2.

CKR2 (rationality): The players are rational in the sense of expected utility theory – they always choose strategies that maximize their individual expected payoffs, relative to their knowledge and beliefs – and this is *common knowledge* in the game.

The concept of *common knowledge* was introduced by Lewis [16] and formalized by Aumann [1]. A proposition is common knowledge if every player knows it to be true, knows that every other player knows it to be true, knows that every other player knows that every other player knows it to be true, and

so on. This is an everyday phenomenon that occurs, for example, whenever a public announcement is made, so that everyone present not only knows it, but knows that others know it, and so on [21].

Key Concepts

Other key concepts of game theory are most easily explained by reference to a specific example. Figure 1 depicts the best known of all strategic games, the *Prisoner’s Dilemma game*. The figure shows its *payoff matrix*, which specifies the game in *normal form* (or *strategic form*), the principal alternative being *extensive form*, which will be illustrated in the section titled Subgame-perfect and Trembling-hand Equilibria. Player I chooses between the rows labeled *C* (cooperate) and *D* (defect), Player II chooses between the columns labeled *C* and *D*, and the pair of numbers in each cell represent the payoffs to Player I and Player II, in that order by convention. In *noncooperative game theory*, which is being outlined here, it is assumed that the players choose their strategies simultaneously, or at least independently, without knowing what the coplayer has chosen. A separate branch of game theory, called *cooperative game theory*, deals with games in which players are free to share the payoffs by negotiating coalitions based on binding and enforceable agreements. The rank order of the payoffs, rather than their absolute values, determines the strategic structure of a game. Replacing the payoffs 5, 3, 1, 0 in Figure 1 by 4, 3, 2, 1, respectively, or by 10, 1, −2, −20, respectively, changes some properties of the game but leaves its strategic structure (Prisoner’s Dilemma) intact.

The Prisoner’s Dilemma game is named after an interpretation suggested in 1950 by Tucker [30] and popularized by Luce and Raiffa [18, pp. 94–97]. Two people, charged with joint involvement in a crime, are held separately by the police, who have insufficient evidence for a conviction unless at least one of them discloses incriminating evidence. The

police offer each prisoner the following deal. If neither discloses incriminating evidence, then both will go free; if both disclose incriminating evidence, then both will receive moderate sentences; and if one discloses incriminating evidence and the other conceals it, then the former will be set free with a reward for helping the police, and the latter will receive a heavy sentence. Each prisoner, therefore, faces a choice between cooperating with the coplayer (concealing the evidence) and defecting (disclosing it). If both cooperate, then the payoffs are good for both (3, 3); if both defect, then the payoffs are worse for both (1, 1); and if one defects while the other cooperates, then the one who defects receives the best possible payoff and the cooperator the worst, with payoffs (5, 0) or (0, 5), depending on who defects.

This interpretation rests on the assumption that the utility numbers shown in the payoff matrix do, in fact, reflect the prisoners’ preferences. Considerations of loyalty and a reluctance to betray a partner-in-crime might reduce the appeal of being the sole defector for some criminals, in which case that outcome might not yield the highest payoff. But the payoff numbers represent von Neumann–Morgenstern utilities, and they are, therefore, assumed to reflect the players’ preferences after taking into account such feelings and everything else affecting their preferences. Many everyday interactive decisions involving cooperation and competition, trust and suspicion, altruism and spite, threats, promises, and commitments turn out, on analysis, to have the strategic structure of the Prisoner’s Dilemma game [7]. An obvious example is price competition between two companies, each seeking to increase its market share by undercutting the other.

How should a rational player act in a Prisoner’s Dilemma game played once? The first point to notice is that *D* is a *best reply* to both of the coplayer’s strategies. A best reply (or best response) to a coplayer’s strategy is a strategy that yields the highest payoff against that particular strategy. It is clear that *D* is a best reply to *C* because it yields a payoff of 5, whereas a *C* reply to *C* yields only 3; and *D* is also a best reply to *D* because it yields 1 rather than 0. In this game, *D* is a best reply to both of the coplayer’s strategies, which means that defection is a best reply whatever the coplayer chooses. In technical terminology, *D* is a *dominant strategy* for both players. A dominant strategy is one that is a best

		II	
		<i>C</i>	<i>D</i>
I	<i>C</i>	3,3	0,5
	<i>D</i>	5,0	1,1

Figure 1 Prisoner’s Dilemma game

reply to *all* the strategies available to the coplayer (or coplayers, if there are several).

Strategic dominance is a decisive argument for defecting in the one-shot Prisoner’s Dilemma game – it is in the rational self-interest of each player to defect, whatever the other player might do. In general, if a game has a dominant strategy, then a rational player will certainly choose it. A dominated strategy, such as *C* in the Prisoner’s Dilemma game, is *inadmissible*, inasmuch as no rational player would choose it. But the Prisoner’s Dilemma game embodies a genuine paradox, because if both players cooperate, then each receives a better payoff (each gets 3) than if both defect (each gets 1).

Nash Equilibrium

The most important “solution concept” of game theory flows directly from best replies. A *Nash equilibrium* (or *equilibrium point* or simply *equilibrium*) is an outcome in which the players’ strategies are best replies to each other. In the Prisoner’s Dilemma game, joint defection is a Nash equilibrium, because *D* is a best reply to *D* for both players, and it is a unique equilibrium, because no other outcome has this property. A Nash equilibrium has strategic stability, because neither player could obtain a better payoff by choosing differently, given the coplayer’s choice, and the players, therefore, have no reason to regret their own choices when the outcome is revealed.

The fundamental theoretical importance of Nash equilibrium rests on the fact that if a game has a uniquely rational solution, then it must be a Nash equilibrium. Von Neumann and Morgenstern [31, pp. 146–148] established this important result via a celebrated *indirect argument*, the most frequently cited version of which was presented later by Luce and Raiffa [18, pp. 63–65]. Informally, by CKR2, the players are expected utility maximizers, and by CKR1, any rational deduction about the game is common knowledge. Taken together, these premises imply that, in a two-person game, if it is uniquely rational for the players to choose particular strategies, then those strategies must be best replies to each other. Each player can anticipate the coplayer’s rationally chosen strategy (by CKR1) and necessarily chooses a best reply to it (by CKR2); and because the strategies are best replies to each other, they are in

Nash equilibrium by definition. A uniquely rational solution must, therefore, be a Nash equilibrium.

The indirect argument also provides a proof that a player cannot solve a game with the techniques of standard (individual) decision theory (*see* strategies of decision making) by assigning subjective probabilities to the coplayer’s strategies as if they were states of nature and then simply maximizing SEU. The proof is by *reductio ad absurdum*. Suppose that a player were to assign subjective probabilities and maximize SEU in the Prisoner’s Dilemma game. The specific probabilities are immaterial, so let us suppose that Player I, for whatever reason, believed that Player II was equally likely to choose *C* or *D*. Then, Player I could compute the SEU of choosing *C* as $1/2(3) + 1/2(0) = 1.5$, and the SEU of choosing *D* as $1/2(5) + 1/2(1) = 3$; therefore, to maximize SEU, Player I would choose *D*. But if that were a rational conclusion, then by CKR1, Player II would anticipate it, and by CKR2, would choose (with certainty) a best reply to *D*, namely *D*. This leads immediately to a contradiction, because it proves that Player II was not equally likely to choose *C* or *D*, as assumed from the outset. The only belief about Player II’s choice that escapes contradiction is that Player II will choose *D* with certainty, because joint defection is the game’s unique Nash equilibrium.

Nash proved in 1950 [22] that every game with a finite number of pure strategies has at least one equilibrium point, provided that the rules of the game allow mixed strategies to be used. The problem with Nash equilibrium as a solution concept is that many games have multiple equilibria that are nonequivalent and noninterchangeable, and this means that game theory is systematically indeterminate. This is illustrated in Figure 2, which shows the payoff matrix of the *Stag Hunt game*, first outlined in 1755 by Rousseau [26, Part II, paragraph 9], introduced into the literature of game theory by Lewis [16, p. 7], brought to prominence by Aumann [1], and discussed

		II	
		<i>C</i>	<i>D</i>
I	<i>C</i>	9,9	0, 8
	<i>D</i>	8,0	7,7

Figure 2 Stag Hunt game

in an influential book by Harsanyi and Selten [13, pp. 357–359]. It is named after Rousseau’s interpretation of it in terms of a hunt in which joint cooperation is required to catch a stag, but each hunter is tempted to go after a hare, which can be caught without the other’s help. If both players defect in this way, then each is slightly less likely to succeed in catching a hare, because they may end up chasing the same one.

This game has no dominant strategies, and the (C, C) and (D, D) outcomes are both Nash equilibria because, for both players, C is the best reply to C , and D is the best reply to D . In fact, there is a third Nash equilibrium – virtually all games have odd numbers of equilibria – in which both players use the mixed strategy $(7/8C, 1/8D)$, yielding expected payoffs of $63/8$ to each. The existence of multiple Nash equilibria means that formal game theory specifies no rational way of playing this game, and other psychological factors are, therefore, likely to affect strategy choices.

Payoff Dominance

Inspired by the Stag Hunt game, and in an explicit attempt to provide a method for choosing among multiple equilibria, Harsanyi and Selten’s *General Theory of Equilibrium Selection in Games* [13] introduced as axioms two auxiliary principles. The first and most important is the *payoff-dominance* principle, not to be confused with strategic dominance, discussed in the section titled Key Concepts. If e and f are two equilibria in a game, then e payoff-dominates (or Pareto-dominates) f if, and only if, e yields a strictly greater payoff to every player than f does. The payoff-dominance principle is the proposition that if one equilibrium payoff-dominates all others in a game, then the players will play their parts in it by choosing its strategies. Harsanyi and Selten suggested that this principle should be regarded as part of every player’s ‘concept of rationality’ and should be common knowledge among the players.

In the Stag Hunt game, (C, C) payoff-dominates (D, D) , and it also payoff-dominates the mixed-strategy equilibrium in which both players choose $(7/8C, 1/8D)$; therefore, the payoff-dominance principle requires both players to choose C . But this assumption requires collective reasoning that goes beyond the orthodox rationality assumption of CKR2.

Furthermore, it is not intuitively obvious that players should choose C , because, by so doing, they risk the worst possible payoff of zero. The D strategy is a far safer choice, risking a worst possible payoff of 7. This leads naturally to Harsanyi and Selten’s secondary criterion of selection among multiple equilibria, called the *risk-dominance* principle, to be used only if payoff dominance fails to yield a determinate solution. If e and f are any two equilibria in a game, then e risk-dominates f if, and only if, the minimum possible payoff resulting from the choice of e is strictly greater than the minimum possible payoff resulting from the choice of f , and players who follow the risk-dominance principle choose its strategies. In the Stag Hunt game, D risk-dominates C for each player, but the payoff-dominance principle takes precedence, because, in this game, it yields a determinate solution.

Subgame-perfect and Trembling-hand Equilibria

Numerous refinements of the Nash equilibrium concept have been suggested to deal with the problem of multiple Nash equilibria and the consequent indeterminacy of game theory. The most influential of these is the *subgame-perfect equilibrium*, introduced by Selten [27]. Selten was the first to notice that some Nash equilibria involve strategy choices that are clearly irrational when examined from a particular point of view. A simple example is shown in Figure 3.

The payoff matrix in Figure 3(a) specifies a game in which both (C, C) and (D, D) are Nash equilibria, but only (C, C) is subgame perfect. The Nash equilibrium (D, D) is not only weak (because 3 is not strictly greater than 3 for Player II) but also imperfect, because it involves an irrational choice from

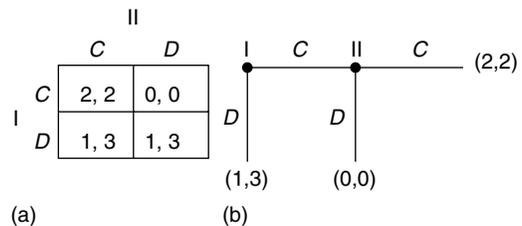


Figure 3 Subgame-perfect equilibrium

Player II. This emerges most clearly from an examination of the *extensive form* of the game, shown in Figure 3(b). The extensive form is a *game tree* depicting the players' moves as if they occurred sequentially. This extensive form is read from Player I's move on the left. If the game were played sequentially, and if the second decision node were reached, then a utility-maximizing Player II would choose *C* at that point, to secure a payoff of 2 rather than zero. At the first decision node, Player I would anticipate Player II's reply, and would, therefore, choose *C* rather than *D*, to secure 2 rather than 1. This form of analysis, reasoning backward from the end, is called *backward induction* and is the basic method of finding subgame-perfect equilibria. In this game, it shows that the (*D*, *D*) equilibrium could not be reached by rational choice in the extensive form, and that means that it is imperfect in the normal form. By definition, a subgame-perfect equilibrium is one that induces payoff-maximizing choices in every branch or subgame of its extensive form.

In a further refinement, Selten [28] introduced the concept of the *trembling-hand equilibrium* to identify and eliminate imperfect equilibria. At every decision node in the extensive form or game tree, there is assumed to be a small probability ε (epsilon) that the player acts irrationally and makes a mistake. The introduction of these error probabilities, generated by a random process, produces a *perturbed game* in which every move that could be played has some positive probability of being played. Assuming that the players' trembling hands are common knowledge in a game, Selten proved that only the subgame-perfect equilibria of the original game remain equilibria in the perturbed game, and they continue to be equilibria as the probability ε tends to zero. According to this widely accepted refinement of the equilibrium concept, the standard game-theoretic rationality assumption (CKR2) is reinterpreted as a limiting case of incomplete rationality.

Experimental Games

Experimental games have been performed since the 1950s in an effort to understand the strategic interaction of human decision makers with bounded rationality and a variety of nonrational influences on their behavior (for detailed reviews, see [6, 11, 15, Chapters 1–4], [17, 25]). Up to the end of the 1970s,

experimental attention focused largely on the Prisoner's Dilemma and closely related games. The rise of behavioral economics in the 1980s led to experiments on a far broader range of games – see [4, 5].

The experimental data show that human decision makers deviate widely from the rational prescriptions of orthodox game theory. This is partly because of bounded rationality and severely limited capacity to carry out indefinitely iterated recursive reasoning ('I think that you think that I think...') (see [8, 14, 29]), and partly for a variety of unrelated reasons, including a strong propensity to cooperate, even when cooperation cannot be justified on purely rational grounds [7].

Evolutionary Game Theory

The basic concepts of game theory can be interpreted as elements of the theory of natural selection as follows. Players correspond to individual organisms, strategies to organisms' genotypes, and payoffs to the changes in their Darwinian fitness – the numbers of offspring resembling themselves that they transmit to future generations. In evolutionary game theory, the players do not choose their strategies rationally, but natural selection mimics rational choice. Maynard Smith and Price [20] introduced the concept of the *evolutionarily stable strategy* (ESS) to handle such games. It is a strategy with the property that if most members of a population adopt it, then no mutant strategy can invade the population by natural selection, and it is, therefore, the strategy that we should expect to see commonly in nature. An ESS is invariably a symmetric Nash equilibrium, but not every symmetric Nash equilibrium is an ESS.

The standard formalization of ESS is as follows [19]. Suppose most members of a population adopt strategy *I*, but a small fraction of mutants or invaders adopt strategy *J*. The expected payoff to an *I* individual against a *J* individual is written $E(I, J)$, and similarly for other combinations strategies. Then, *I* is an ESS if *either* of the conditions (1) or (2) below is satisfied:

$$E(I, I) > E(J, I) \tag{1}$$

$$E(I, I) = E(J, I), \text{ and } E(I, J) > E(J, J) \tag{2}$$

Condition (1) or (2) ensures that J cannot spread through the population by natural selection. In addition, differential and difference equations called *replicator dynamics* have been developed to model the evolution of a population under competitive selection pressures. If a population contains k genetically distinct types, each associated with a different pure strategy, and if their proportions at time t are $x(t) = (x_1(t), \dots, x_k(t))$, then the replicator dynamic equation specifies the population change from $x(t)$ to $x(t + 1)$.

Evolutionary game theory turned out to solve several long-standing problems in biology, and it was described by Dawkins as ‘one of the most important advances in evolutionary theory since Darwin’ [9, p. 90]. In particular, it helped to explain the evolution of cooperation and altruistic behavior – conventional (ritualized) rather than escalated fighting in numerous species, alarm calls by birds, distraction displays by ground-nesting birds, and so on.

Evolutionary game theory is also used to study adaptive learning in games repeated many times. Evolutionary processes in games have been studied analytically and computationally, sometimes by running simulations in which strategies are pitted against one another and transmit copies of themselves to future generations in proportion to their payoffs (see [2, 3, Chapters 1, 2; 23, 24]).

References

- [1] Aumann, R.J. (1976). Agreeing to disagree, *Annals of Statistics* **4**, 1236–1239.
- [2] Axelrod, R. (1984). *The Evolution of Cooperation*, Basic Books, New York.
- [3] Axelrod, R. (1997). *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*, Princeton University Press, Princeton.
- [4] Camerer, C.F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, Princeton.
- [5] Camerer, C.F., Loewenstein, G. & Rabin, M., eds (2004). *Advances in Behavioral Economics*, Princeton University Press, Princeton.
- [6] Colman, A.M. (1995). *Game Theory and its Applications in the Social and Biological Sciences*, 2nd Edition, Routledge, London.
- [7] Colman, A.M. (2003a). Cooperation, psychological game theory, and limitations of rationality in social interaction, *The Behavioral and Brain Sciences* **26**, 139–153.
- [8] Colman, A.M. (2003b). Depth of strategic reasoning in games, *Trends in Cognitive Sciences* **7**, 2–4.
- [9] Dawkins, R. (1976). *The Selfish Gene*, Oxford University Press, Oxford.
- [10] Dimand, M.A. & Dimand, R.W. (1996). *A History of Game Theory (Vol 1): From the Beginnings to 1945*, Routledge, London.
- [11] Foddy, M., Smithson, M., Schneider, S. & Hogg, M., eds (1999). *Resolving Social Dilemmas: Dynamic, Structural, and Intergroup Aspects*, Psychology Press, Hove.
- [12] Harsanyi, J.C. (1967–1968). Games with incomplete information played by “Bayesian” players, Parts I–III, *Management Science* **14**, 159–182, 320–334, 486–502.
- [13] Harsanyi, J.C. & Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge.
- [14] Hedden, T. & Zhang, J. (2002). What do you think I think you think? Strategic reasoning in matrix games, *Cognition* **85**, 1–36.
- [15] Kagel, J.H. & Roth, A.E., eds (1995). *Handbook of Experimental Economics*, Princeton University Press, Princeton.
- [16] Lewis, D.K. (1969). *Convention: A Philosophical Study*, Harvard University Press, Cambridge.
- [17] Liebrand, W.B.G., Messick, D.M. & Wilke, H.A.M., eds (1992). *Social Dilemmas: Theoretical Issues and Research Findings*, Pergamon, Oxford.
- [18] Luce, R.D. & Raiffa, H. (1957). *Games and Decisions: Introduction and Critical Survey*, Wiley, New York.
- [19] Maynard Smith, J. (1982). *Evolution and the Theory of Games*, Cambridge University Press, Cambridge.
- [20] Maynard Smith, J. & Price, G.R. (1973). The logic of animal conflict, *Nature* **246**, 15–18.
- [21] Milgrom, P. (1981). An axiomatic characterization of common knowledge, *Econometrica* **49**, 219–222.
- [22] Nash, J.F. (1950). Equilibrium points in n -person games, *Proceedings of the National Academy of Sciences of the United States of America* **36**, 48–49.
- [23] Nowak, M.A., May, R.M. & Sigmund, K. (1995). The arithmetics of mutual help, *Scientific American* **272**(6), 76–81.
- [24] Nowak, M.A. & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game, *Nature* **364**, 56–58.
- [25] Pruitt, D.G. & Kimmel, M.J. (1977). Twenty years of experimental gaming: critique, synthesis, and suggestions for the future, *Annual Review of Psychology* **28**, 363–392.
- [26] Rousseau, J.-J. (1755). Discours sur l’origine d’inegalité parmi les hommes [discourse on the origin of inequality among men], in *Oeuvres Complètes*, Vol. 3, J.-J. Rousseau, Edition Pléiade, Paris.
- [27] Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit [Game-theoretic treatment of an oligopoly model with demand inertia], *Zeitschrift für die Gesamte Staatswissenschaft* **121**, 301–324, 667–689.
- [28] Selten, R. (1975). Re-examination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory* **4**, 25–55.

- [29] Stahl, D.O. & Wilson, P.W. (1995). On players' models of other players: theory and experimental evidence, *Games and Economic Behavior* **10**, 218–254.
- [30] Tucker, A. (1950/2001). A two-person dilemma, in *Readings in Games and Information*, E. Rasmusen, ed., Blackwell, Oxford, pp. 7–8 (Reprinted from a handout distributed at Stanford University in May 1950).
- [31] von Neumann, J. & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*, 2nd Edition, 1947; 3rd Edition, 1953, Princeton University Press, Princeton.

ANDREW M. COLMAN