The Counterintuitive in Conflict and Cooperation

Joel E. Cohen

Conflicts and cooperation offer two kinds of surprises: those generated by a participant and those generated by the situation. Surprises of the first kind are familiar. For example, in *The Art of War* (1521), the Florentine bureaucrat Niccolò Machiavelli counseled military commanders to surprise their opponents; he gave many historical examples of such surprises. How best to achieve surprise in political and military conflicts is still under analysis (e.g., Axelrod 1979).

A second kind of surprise arises when, because of the structure of interaction among opponents or collaborators, intuitively reasonable actions or policies lead to an outcome that surprises everybody. I will call such surprises structural surprises. Examples could be drawn from historical case studies (e.g., Tuchman 1962; Bracken 1983, 1988; Carter et al. 1987). However, any example is open to the argument that the participants did not analyze the situation sufficiently.

Mathematical models of real-life problems suggest that following one's intuition can lead to results that nobody wants

Structural surprises arise even in well defined, highly simplified mathematical models. As examples, the iterated prisoners' dilemma (Axelrod 1984; Downs et al. 1985), rules for voting (Brams and Fishburn 1983; Saari 1987b), social choice (Arrow 1963; Saari 1987a), and proportional representation (Balinski and Young 1982) all sometimes generate surprising outcomes from apparently reasonable rules of thumb. Mathematics itself is rich in paradoxes and surprises (e.g., Maxwell 1959; Bunch 1982; Székely 1987), so it is unsurprising that concrete interpretations of mathematics can be surprising.

In this article, recent examples of structural surprise will be drawn from idealized models of conflict and cooperation. The models describe communication networks under attack, pursuit and evasion, negotiations for consensus, and congested traffic networks. While these simplified models appear transparent to the intuition, mathematical analysis shows that the use of reasonable and traditional rules of thumb can lead to surprising results for all participants. Mathematical analyses of these models appear elsewhere. Here I describe only the models and their properties.

Because the training and intuition of different people are different, one person's surprise is often another person's so-what. In the examples that follow, the claims that findings are surprising are accompanied by efforts to build up the appropriate intuition. To each such claim, the qualification applies that the finding is surprising at least to some people.

It might be inferred from these examples that real conflict and cooperation are, or can be, so complex and counterintuitive that they should be left to the management of experts, particularly military experts. Such an inference is not justified. On the contrary, these examples show that conflict and cooperation require the fresh and unprejudiced analysis of thinkers who have no commitment, historical or bureaucratic, to established rules of thumb.

When does redundancy enhance reliability?

Figure 1 is a sketch of the United States strategic command and control system (Bracken 1983). Presumably an analogous network of command and control exists for the nuclear forces of the Soviet Union (Steinbruner 1984; Zraket 1984).

In any strategic conflict, each side has an obvious interest in preserving the integrity of its own system of command and communication. In the face of an impending atomic attack, the vulnerability of one side's command system "presents a much more powerful incentive to initiate attack before damage has actually been suffered, an incentive that is driven . . . by practical fears of a decisive defeat in a war that cannot be avoided" (Steinbruner 1984).

Each side also has a stake in the integrity of the communication network of the opposing side. If, after a nuclear attack, the command network of an opponent were no longer connected (in the sense that each commander of nuclear forces could communicate with every other), the commander of an isolated component might react to the attack without knowing the responses or intentions of other commanders. Even if the National

Photographs on the facing page illustrate classes of problems discussed in this article. ***Clockwise, from upper left:*** **reliability of command systems against nuclear attack (an IM-99 Bomarc interceptor missile takes off from Patrick Air Force Base); pursuit and evasion (the Keystone cops); approaching consensus through negotiations (Mikhail Gorbachev and Ronald Reagan at a summit meeting in Geneva, 21 November 1985); traffic flow on congested roads (Interstate 80 in San Francisco). All photos from the Bettmann Archive.**

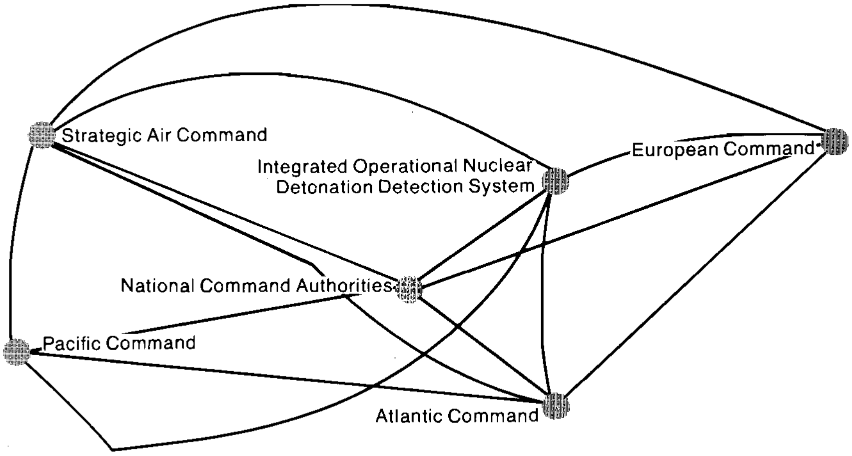

Figure 1. A stylized command system for the United States strategic forces shows redundant lines of communication. A conventional rule of thumb is that redundancy enhances reliability. The question is, when does this rule hold? (After Bracken 1983.)

Command Authorities (NCA) wished to limit further hostilities, they would have lost control of some forces, say, the Pacific Command (PAC). Fearing that PAC may continue to attack and that the Soviet Union may react to the PAC attacks by attacking the NCA or the rest of the United States, the NCA may see no incentive to limit further hostilities (Bracken 1983).

These arguments and others indicate that a potential or actual loss of connectivity in the command system makes the prevention and termination of conflict highly unstable (Kettelle 1981).

A communication network can be modeled as a graph, and a network under attack or with other sources of failure as a random graph (Bollobás 1985; Palmer 1985). A graph, in the sense used here, consists of vertices (or points) and edges. Each vertex corresponds to a command center or other node in a communication network. Each edge represents a two-way communication link between two command centers. In a random graph, an edge may fail, leaving intact the vertices connected by the edge, or a vertex may fail, destroying all the edges connected to the vertex.

As a first approximation, simple enough to analyze mathematically and complex enough to shed light on reality, suppose that the vertices are hardened against attack, at least by comparison with the edges, and that only the edges fail. Assume that the network has n fixed vertices. In Figure 1, $n = 6$. In the still simpler hypothetical random graph in Figure 2, $n = 3$. Assume that the edge between vertex i and vertex j will be present and operational postattack with probability p_{ij} and will be not operational (failed) with probability $1 - p_{ij}$. Assume that each edge works or fails postattack independently of every other edge. These assumptions regarding fixed vertices and independently randomly failing edges define an anisotropic edge-random graph. The graph is called anisotropic because the probabilities for failure may be different for different edges.

The graph, pre- or postattack, is said to be connected if at the time of observation, for every pair of vertices i and j, there is a path of edges between i and j. The smallest number of edges required to connect a graph with n vertices is $n - 1$. In any connected graph with n vertices and $n - 1$ edges, there is exactly one path from any vertex i to any other vertex j. A spanning tree of a graph is defined to be a subgraph of $n - 1$ edges that connects all n vertices. For example, in Figure 2, one spanning tree is the pair of edges from vertex 1 to vertex 2 and from vertex 1 to vertex 3: these two edges connect all three vertices of the graph. A graph is connected if and only if it has at least one spanning tree as a subgraph.

The reliability of a random graph is defined to be the probability that it is connected. The reliability, denoted by r, is thus the probability that the random graph has one or more spanning trees.

The design of a reliable network would be easy if cost were no object. The designer would simply maximize the probability of survival of each edge as far as physically possible. For real networks, the design problem is to maximize reliability given limited resources. A conventional rule of thumb in the design of networks is that, subject to a given total cost, redundancy enhances reliability. To find out when this rule of thumb is valid, one needs a precise measure of the redundancy of an edge-random graph, as well as a measure of cost.

The average or expected number of spanning trees postattack, denoted by t (for trees), measures the average number of different ways the network could be connected after edges have failed. This seems a reasonable measure of redundancy, useful as a first approximation. To get a feeling for how the redundancy t is computed, consider the simple random graph in Figure 2. There is a certain probability that all three edges will fail postattack; in this case, there will be no spanning trees. Likewise, if any pair of edges fails, there will be no spanning trees. If only one edge fails, there will be exactly one spanning tree left, consisting of the two surviving edges. Finally, there is a certain probability that no edges will fail postattack; in this case, there will be three spanning trees, each possible pair of edges being one spanning tree. The redundancy t is just the average of the number of spanning trees in each of these cases, weighted by the corresponding probability of each case.

The general formulas for the reliability r and the redundancy t of random graphs with an arbitrary number n of vertices are complex. But the formulas when $n =$ 3 (as in Fig. 2) are simple:

$$r = p_{12}p_{13} + p_{12}p_{23} + p_{13}p_{23} - 2p_{12}p_{13}p_{23}$$
$$t = p_{12}p_{13} + p_{12}p_{23} + p_{13}p_{23}.$$

The rule of thumb that, for a fixed cost, redundancy enhances reliability can be translated into a testable hypothesis, assuming that the formulas for r and t capture what designers mean by reliability and redundancy, respectively. To state this hypothesis concisely, we need a name for the list of edge probabilities of a

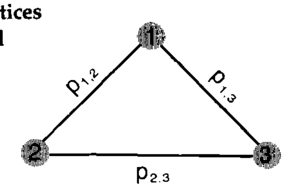

Figure 2. A graph with three vertices reduces Figure 1 to mathematical essentials. The vertices are analogous to command centers, and the edges (which connect vertices) are analogous to lines of communication. The graph is "random" in the sense that any edge may randomly fail. The probability that an edge will survive an attack is denoted here as p (for example, $p_{1,2}$ refers to the edge between vertices 1 and 2). In Figures 3 and 4, these probabilities are given as specific numbers.

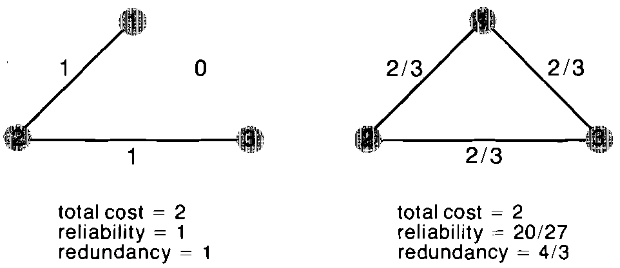

Figure 3. Two random graphs show different strategies for surviving an attack. Both have the same total cost (computed using the function c_1 described in the text). In the graph on the left, two edges are hardened to guarantee survival ($p = 1$) while the third edge is sacrificed ($p = 0$). Each vertex is thus assured of communication with the others (1 and 3 can communicate through 2). In the graph on the right, each edge is hardened to give it a 2/3 probability of surviving. The graph on the left is more reliable even though the one on the right has more redundancy built into the network, according to a mathematical model described in the text.

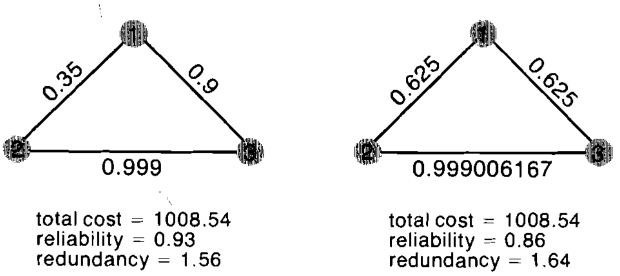

Figure 4. Two random graphs have the same total cost (calculated using the function c_2 described in the text), but their edges have different probabilities of survival. The graph on the right is less reliable even though more redundant.

random graph. Let us call that list P; for example, in Figure 2, $P = (p_{12}, p_{13}, p_{23})$. We also need a name for the total cost of attaining the probabilities of edge survival P; let us call that total cost $c(P)$. Now suppose P_1 and P_2 are edge-probability lists of two different random graphs that have the same number n of vertices. Suppose that $c(P_1) = c(P_2)$, so that achieving the survival probabilities of the edges in the two random graphs costs the same total amount. If redundancy enhances reliability, then it should be true that, whenever $t(P_1) < t(P_2)$, it is also true that $r(P_1) \leq r(P_2)$. In words, the rule of thumb suggests that, for a fixed total cost, if the network described by P_1 is less redundant than the network described by P_2, then the network described by P_1 should be less reliable than the network described by P_2.

One measure of the total cost associated with an edge-probability matrix P is the sum of all the edge probabilities:

$$c_1(P) = \Sigma_{i<j}\, p_{ij}.$$

The designer faces the choice of allocating a fixed sum of probabilities among the possible edges.

Figure 3 contradicts the expectation that, for a given total cost, redundancy enhances reliability (Cohen 1986). In the graph on the left, the edge between vertices 1 and 2 is guaranteed not to fail (has probability 1) and similarly for the edge between vertices 2 and 3. Thus the network has reliability 1; it is perfectly reliable. The total cost of the network, using c_1, is 2. In the graph on the right, the probabilities of all three possible edges are 2/3. The sum of edge probabilities (total cost) is the same as in the graph on the left. Though the random graph on the right has greater redundancy than that on the left, it has lower reliability. Hence, redundancy need not enhance reliability.

In this design problem, when the total cost $c_1(P)$ exceeds 1, the designer apparently can guarantee that at least one edge will work with certainty, that is, with probability 1. This seems unrealistic. A better measure of cost is the sum, over all edges, of some function of edge probabilities that is 0 when the edge probability is 0, increases faster and faster as the edge probability increases (is convex in the edge probability), and goes to infinity as the edge probability approaches 1. For example:

$$c_2(P) = \Sigma_{i<j}\, p_{ij}/(1 - p_{ij}).$$

Figure 4 shows two random networks with the same total cost using c_2. In the graph on the right, the edge probabilities between vertices 1 and 2 and between vertices 1 and 3 are the average of the two corresponding edge probabilities in the graph on the left. The edge probability between vertices 2 and 3 was chosen to make the total cost, using c_2, of the network on the right equal to that of the network on the left. The redundancy of the network on the right exceeds that of the one on the left, but the reliability is decreased.

These examples raise a question for future research: When does redundancy enhance reliability?

Pursuit-evasion games on graphs

Many physical conflicts include pursuit and evasion, hunting and fleeing, chase and flight. The mathematical theory of pursuit and evasion, in its minimally realistic forms, is notoriously difficult and resistant to the intuition. An extremely simplified, admittedly unrealistic, formulation of pursuit and evasion offered here combines the theory of simple graphs with von Neumann's classical theory of zero-sum two-person games (von Neumann 1928; Dresher 1961).

Suppose that the field of combat for pursuit and evasion can be modeled by a connected simple graph G. G consists of a finite number of vertices, which may be thought of as physical locations, and some edges. An edge between two vertices models a possible route of travel in either direction between the two corresponding locations. (Loops and multiple edges are excluded. A loop is an edge from a vertex to itself. Multiple edges are more than one edge between a given pair of vertices.) That the graph is connected means, as before, that there exists a path or sequence of edges from any vertex to any other.

Suppose there are two players: Red, the pursuer, and Blue, the evader. Red and Blue each choose a vertex independently and without knowing the choice of the other player. If Red chooses vertex i and Blue chooses vertex j, then Red must pay Blue an amount in dollars equal to the distance $d(i,j)$ between vertices i and j. The distance $d(i,j)$ is defined as the smallest number of edges that must be traversed to travel from i to j (or from j to i). For example, in the graph in Figure 5, the path from vertex 1 to vertex 2 that passes through vertex 3 traverses 4 edges, while the path from vertex 1 to vertex 2 that does not pass through vertex 3 traverses 3 edges; therefore, the distance from vertex 1 to vertex 2 is 3, the lesser of the two lengths. In this game, the pursuer Red has an incentive to choose a vertex as close as possible to the vertex chosen by the evader Blue, and Red's loss is Blue's gain.

A pure strategy in this game is a choice of a single vertex (with probability 1). An optimal strategy for Red is one that minimizes the maximum Red must pay Blue, whatever Blue does. Red's optimal pure strategies are easy to describe, once some standard graph-theoretic terms (Harary 1969) are defined and illustrated.

The eccentricity of a vertex i is the maximum distance from i of any vertex j in the graph. For example, in Figure 5, the eccentricity of vertex 1 is 3 because vertex 1 is 3 edges distant from vertex 2 and every other vertex is closer to vertex 1. Similarly, the eccentricities of vertices 3, 4, and 5 are all 2. The radius of a graph is the smallest eccentricity of any vertex. In Figure 5, the radius of the graph is 2 because some vertices have eccentricity 2 and none have smaller eccentricity. The center of a graph is the set of all vertices whose eccentricity equals the radius of the graph. In Figure 5, the center of the graph consists of the vertices 3, 4, and 5.

If Red chooses any vertex in the center of the graph, that position minimizes the maximum Red must pay Blue, no matter which vertex Blue chooses. Hence, the center of the graph contains all the optimal pure strategies for Red. Intuitively speaking, Red can minimize the farthest Blue can run from him by sitting somewhere in the center of the graph.

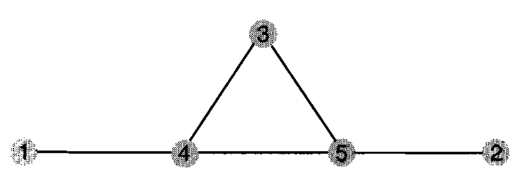

Figure 5. In graphs of pursuit and evasion, vertices are analogous to stopping points and edges to routes of travel. One way of using this graph to simulate real conflicts is to have the pursuer and evader each choose a vertex independently, without knowing the other's choice. The pursuer then pays the evader an amount of dollars equal to the distance between them. A pursuer might intuitively choose vertex 3, which minimizes the maximal distance to any vertex. However, if the evader chooses vertices at random, every optimal strategy for the pursuer avoids vertex 3 completely.

Figure 6. In this graph of 14 vertices, a pursuer might intuitively choose to wait at the center (vertex 1), but no optimal strategy includes this vertex. The pursuer would do better to wait at vertices 2, 3, or 4.

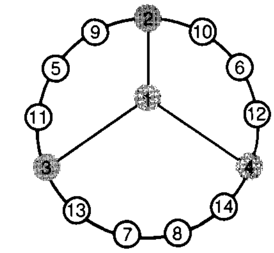

However, Red and Blue can play using pure strategies or mixed strategies. If Red randomizes his or her choice of a vertex, say by flipping a coin, spinning a roulette wheel, or using any other random device, or if Red plays the pursuit-evasion game repeatedly and randomly chooses different vertices in different plays of the game, Red is using not a pure strategy but a mixed strategy.

The easy analysis of pure strategies leads to a plausible rule of thumb for Red: in optimal mixed strategies, Red should assign a positive probability to some or all of the vertices in the center of the graph and zero probabilities to the vertices outside the center of the graph. It also seems plausible that for every vertex in the center of a graph, there exists at least one mixed strategy that assigns positive probability to that vertex.

These plausible rules of thumb are wrong. In the graph in Figure 5, vertex 3, which belongs to the center, is assigned no probability by any optimal mixed strategy for Red (Chung et al. 1987). To see why, suppose that Red did assign some positive probability, say ε, to vertex 3, and that Blue chose vertex 1 with probability 1/2 and vertex 2 with probability 1/2. Then Red must assign its remaining probability $1 - \varepsilon$ to the remaining vertices 1, 2, 4, and 5. The distance from any one of these four vertices to vertex 1, plus the distance from that same vertex to vertex 2, is exactly 3, so the average cost to Red of probability assigned to any of these four vertices is just 3/2. The sum of the distances from vertex 3 to vertices 1 and 2, however, is 4, so the average cost to Red of the probability ε assigned to vertex 3 is $(4/2)\varepsilon > 0$. Thus when Red assigns positive probability to vertex 3, Red's average cost strictly exceeds 3/2. Red can do better by sitting with probability 1/2 on each of the vertices 4 and 5. The latter strategy assures that Red will not have to pay more than 3/2 on average.

In the graph in Figure 6, no optimal mixed strategy for Red places any probability in the center of the graph, which here consists just of vertex 1. A short mathematical argument (Chung et al. 1987) shows that Red can improve on any strategy that uses vertex 1 with positive probability by a strategy that uses vertex 2 with probability 1/5 and vertices 3 and 4 each with probability 2/5. (The lack of symmetry in these probabilities reflects the lack of symmetry in the graph.) Intuitively, in this specially contrived example, vertices 2, 3, and 4 are closer to where the action is, though they all fall outside the center of the graph.

The examples in Figures 5 and 6 show that a plausible rule of thumb based on pure strategies can be misleading when applied to mixed strategies. What is optimal when players are constrained to make a single choice with probability 1 may be far from optimal when players' actions are uncertain or contain a random element.

Approaching consensus can be delicate

Suppose that two negotiators (call them Red and Blue again) are trying to reach agreement on an estimate of some unknown numerical quantity F (for example, the forces permitted under an arms agreement). A widely studied model (French 1956; Harary 1959; Roberts 1976; Genest and Zidek 1986) of an iterative negotiation supposes that at the start of the negotiation Red proposes the estimate F_R^0 and Blue proposes an estimate F_B^0. After completing k stages of negotiation (initially $k = 0$), Red forms a new estimate F_R^{k+1} of F by taking a weighted average of his own most recent estimate F_R^k and Blue's most recent estimate F_B^k, using weights that may depend on the stage of the negotiation (DeGroot 1974). Thus $F_R^{k+1} = a^{(k)}F_R^k + (1 - a^{(k)})F_B^k$, where $0 \leq a^{(k)} \leq 1$. Blue similarly calculates his new estimate F_B^{k+1} as a weighted sum of his own and Red's most recent estimates, using his own weights, which may also depend on the stage of the negotiation. For simplicity, let us suppose here that the weights Blue uses at each stage are symmetrical to the weights Red uses at each stage, meaning for example that if Red weights his own most recent estimate by 3/4 and weights Blue's most recent estimate by 1/4 then Blue weights Blue's own most recent estimate by 3/4 and Red's most recent estimate by 1/4. In general, suppose that $F_B^{k+1} = (1 - a^{(k)})F_R^k + a^{(k)}F_B^k$. The weights used by both negotiators at stage k are described by $a^{(k)}$.

The negotiators approach consensus if, in the limit as the number k of stages in the negotiation becomes large, the absolute difference $|F_R^k - F_B^k|$ between the estimate F_R^k of Red and the estimate F_B^k of Blue approaches 0, no matter what the initial estimates F^0 of the negotiators may be. In reality, no negotiation can or will continue forever. The definition of approaching consensus just offered remains useful, however, because it discriminates between negotiations in which, if the negotiators choose to continue, they can come arbitrarily close together, and negotiations in which, even if the negotiators persist, they cannot come arbitrarily close together.

If the negotiators constantly use weights of 1/2 for their own and the other negotiator's estimates, they will approach consensus because any initial difference between the estimates of Red and Blue is eliminated at the first stage of the negotiation.

On the other hand, if Red attaches weight 1 to Red's own most recent estimate ($a^{(k)} = 1$) and 0 to Blue's most recent estimate, while Blue attaches weight 0 to Red's most recent estimate and 1 to Blue's own most recent estimate, it is equally clear that the negotiators will never reach consensus. Each negotiator listens only to himself, and the estimates cease to change.

It might be inferred from these examples that real conflict and cooperation should be left to the management of experts. Such an inference is not justified

The negotiators harden their positions if the weight $a^{(k)}$ approaches 1 as k gets large. Surprisingly, the negotiators can come arbitrarily close to agreement in their estimates even if they harden their positions, provided that they do not harden their positions too quickly.

To illustrate this claim, consider, for every stage k of the negotiation, $k = 0, 1, 2, \ldots$, two possible weights $A_k = 1 - 1/(k + 3)$ and $B_k = 1 - 1/(k + 3)^2$. These weights differ only in the exponent of $k + 3$. For large k, A_k and B_k both approach the weight 1. When $a^{(k)} = A_k$, the negotiators harden their positions slowly. When $a^{(k)} = B_k$, the negotiators harden their positions more rapidly.

If $a^{(k)} = A_k$ at every stage k, the negotiators will inevitably reach consensus: whatever the negotiators' initial estimates F^0, the differences between their estimates will vanish as k gets large (Fig. 7). On the other hand, if $a^{(k)} = B_k$ at every stage k, and if the initial estimates of the two negotiators differ, then their subsequent estimates will never become arbitrarily close, though the negotiations persist through any number k of stages (Isaacson and Madsen 1976). It is not at first intuitively obvious why the former weights should lead to consensus while the latter weights should not. One can quickly get a feeling that the two sets of weights are qualitatively different by recalling that the sum for all positive integers k of $1/k$ goes to infinity while the sum

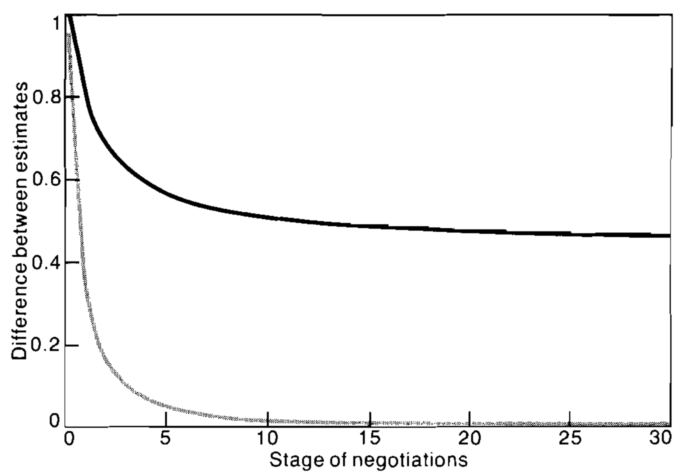

Figure 7. Negotiations approach consensus or deadlock depending on how each side hardens its position. Hardening can be defined mathematically as a "weight" each side attaches to its own and its opponent's latest estimates (or offers). When negotiators harden their positions slowly (*gray*), they approach consensus (the difference between their estimates approaches zero). In this example, the weight that one side attaches to the estimate of the other varies inversely as the stage of the negotiation. When negotiators harden their positions rapidly (*black*), the difference between their estimates declines to 0.44 but never approaches zero. In this example, negotiators vary the weight attached to their opponents' estimates inversely as the square of the stage of negotiation. When randomness affects the weights assigned at different stages, as in real life, small changes in the negotiators' behavior can lead to dramatically different outcomes.

over k of $1/k^2$ is finite. This elementary fact is key to proving that the first set of weights leads to consensus while the second does not.

To make the model of negotiation slightly more realistic, assume that the weights $a^{(k)}$ are not fixed but are affected by random events in the environment of the negotiations. As a first example, suppose that $a^{(k)} = A_k$ with probability p and $a^{(k)} = B_k$ with probability $1 - p$, where $0 \leq p \leq 1$, and that the weights are chosen independently at each stage k. Since the current estimates F^k of the negotiators are a continuous function of both the initial estimates F^0 and of all the weights $a^{(k)}$, it seems plausible that small changes, even over a prolonged period of time, in the behavior of the negotiators should make a small difference in whether the negotiators approach consensus. This rule of thumb suggests that a small change in the probability p of a slow hardening of positions (described by A_k) should cause a small change in the probability that the negotiators reach consensus. This reasonable expectation turns out to be true for a significant range of values of p, but also significantly false. When $p = 0$, that is, when $a^{(k)} = B_k$ with probability 1, the negotiators do not approach consensus. However, for any positive value of p, no matter how close to 0, the negotiators will approach consensus with probability 1, that is, almost surely. Thus the probability that negotiations will approach consensus changes discontinuously from 0 to 1 as p, the probability that the negotiators harden their positions at a slow rate, increases infinitesimally from 0 to any positive value, no matter how small (Cohen et al. 1986).

This first example assumes that the probability that $a^{(k)} = A_k$ is independent of the stage k of the negotiation. Suppose now, as a second example, that the probability that $a^{(k)} = A_k$ is p_k, where p_k is a gradually decreasing function of k. According to this assumption, the negotiators have a gradually declining probability of hardening their own positions at the slow rate described by A_k, and a gradually increasing probability of hardening their positions at the faster rate described by B_k. For purposes of illustration, suppose that $p_k = C/(\ln k)^C$, for some number $C \geq 0$, where ln means the natural logarithm (Cohen et al. 1986).

Because ln k increases quite slowly as k increases, p_k decreases quite slowly as k increases. For values of C less than or equal to 1, the negotiators will approach consensus with probability 1. If C exceeds 1, however, the negotiators will approach consensus with probability 0. Thus an infinitesimal change in the value of C, from less than 1 to greater than 1, destroys the approach to consensus, even though the probabilities p_k appear hardly to have changed at all. This example shows dramatically that small changes in the behavior of negotiators can lead to big changes in the outcome of a negotiation.

Traffic flow on congested roads

Consider drivers who seek to travel by car from an origin to a destination by means of a network of roads. (Think, for local color, of Manhattan, with cars traveling from Battery Park at the south end of the island to George Washington Bridge at the north end, by routes along the west or east sides.) The complexities of a real road network may be approximated by a directed graph (Robinson and Foulds 1980), which consists of vertices and arcs (directed edges). Vertices represent places or junctions of roads. Whenever there is a path from one vertex to another, place an arc (an edge directed from one vertex to another along the direction of possible traffic flow). A directed graph differs from an undirected graph in that each edge is assigned a direction and called an arc. The graphs in Figure 8 are two examples of directed graphs as models of road networks. (The graph on the left represents major arteries along the east and west sides of Manhattan; the graph on the right represents, in addition, an unfortunately nonexistent major artery across the middle of the island.)

The traffic in a road network is described by the number f_{ij} of cars moving along each arc (i,j) from vertex i to vertex j. The cost to each car of travel from i to j along arc (i,j) when the flow is f_{ij} is described by $c_{ij}(f_{ij})$. The cost may be measured by time spent in travel or by gasoline consumed, for example. It is assumed that the cost of

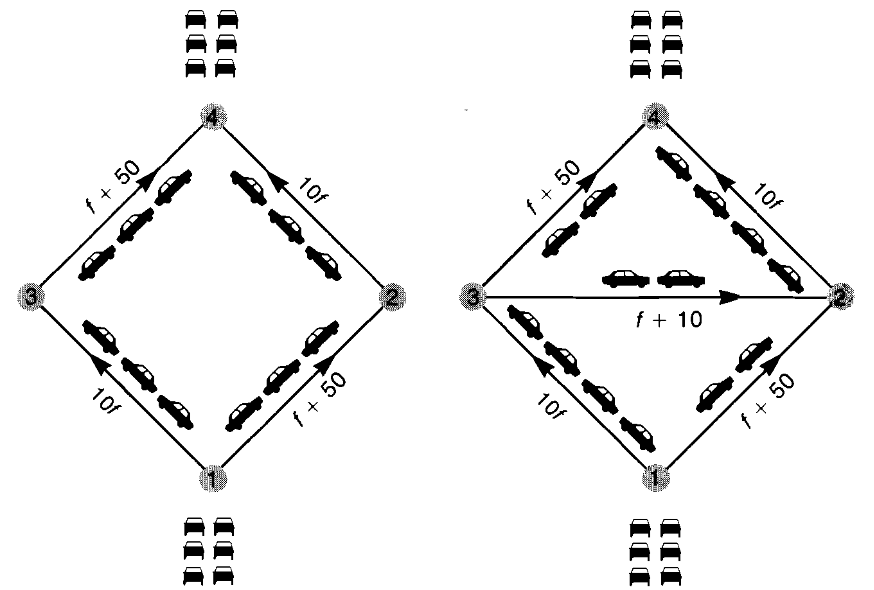

Figure 8. In a traffic network (*left*), cars traverse a grid, from vertex 1 to vertex 4, by way of either vertex 2 or vertex 3. In this example, both routes are congested, and traffic engineers consider building a new route between vertices 2 and 3 (*right*). Will the added flexibility improve circulation? Not for this city. When traffic flow reaches a new equilibrium, the results are worse for everyone: it takes longer to traverse the grid than before. In this figure, the small cars associated with each route show the number of cars taking that route; in both cases, six cars enter and exit the network. The formulas associated with each route give the cost per car (say, in minutes) of traveling a route as a function of the number of cars, *f*, taking that route; for example, when 3 cars travel along a route with cost function 10*f*, each one takes 30 minutes. (After Steinberg and Zangwill 1983.)

traveling along an arc at a given level of traffic is the same for all cars traversing that arc. The cost defined here refers to the cost for each car, not to the aggregate cost for all the cars flowing along the arc.

An arc from i to j is said to be uncongested if one more car in the flow along the arc results in no increase in the cost per car traveling along the arc. A network is said to be uncongested if, at the current levels of flow in each arc, each arc is uncongested.

An arc is said to be congested if one more car flowing along the arc results in an increase in the cost per car traveling along that arc. Each arc in the networks in Figure 8 is congested. The network is said to be congested if one or more of its arcs is congested. Thus both networks in Figure 8 are congested. An arc or network may be uncongested at one level of traffic flow and congested at another.

The total travel cost of a car is the sum of the arc costs over all arcs in the travel route from origin to destination. Each driver is assumed to seek a route from origin to destination that minimizes his or her total travel cost. At equilibrium, the cars are distributed over arcs so that a change of path by any one car would raise the total travel cost for that car, given that all the other cars did not change their routes. At equilibrium, all cars travel by paths that have the same total travel cost. When the network is uncongested, the choice of route by one driver has no effect on the choice of route by another. All drivers travel by minimal cost routes, of which there may be several. In a congested network, on the other hand, a change of path by one car may affect the total travel cost of that car and of many other cars.

In an uncongested network, adding additional arcs cannot increase the total travel cost of any car at equilibrium. An additional arc simply opens more choices to each driver and may lead to no change or to decreases in total travel costs. This fact leads to a tempting rule of thumb for congested networks: more roads in a road network should save time for everybody (or at least somebody).

A startling counterexample to this rule of thumb was discovered by Braess (1968) and is known as Braess's paradox. The network on the left of Figure 8 differs from that on the right only by the addition of an arc from vertex 3 to vertex 2.

Let me describe a route through the network by the vertices through which a car passes. Consider a driver who drives the route (1,3,4) in the network on the left. After the addition of the new arc, if all of the other drivers retain their previous routes, one of the three drivers who originally travels (1,3,4) is in a position to lower his cost by traveling from vertex 3 to vertex 2 via the new arc (incurring a cost of $11 = 1 + 10$ for that arc) and from vertex 2 to vertex 4, incurring a cost of 40 for that arc; his travel cost of 30 for the arc (1,3) remains the same, so his total travel cost would be reduced from $30 + 53 = 83$ to $30 + 11 + 40 = 81$. Therefore he switches from the route (1,3,4) to the route (1,3,2,4). Now there are four cars flowing along the arc (2,4). The three original drivers along the route (1,2,4) are now paying 93 instead of the original 83. One of these three can lower her total travel cost by shifting the first portion of her route from (1,2) to (1,3). If she then continues along the arc (3,4), her total travel cost, 93, will be higher than if she follows the new arc along the route (1,3,2,4), for a total travel cost of 92. At this point, the traffic flows are as shown in the network on the right. Every driver is paying 92. Every driver now finds that he or she would be worse off if he or she changed routes, given what everyone else is doing. Thus when each driver reacts to the reactions of every other driver, the flows and costs stabilize at the new equilibrium shown on the right. The remarkable, and utterly counterintuitive, feature of this new equilibrium is that everyone is worse off! The cost per car has risen from 83 to 92.

Such counterintuitive effects should be by no means unusual in congested road networks (Steinberg and Zangwill 1983). In this and similar examples (Cohen and Kelly, unpubl.), Adam Smith's Invisible Hand leads everyone astray.

Let your imagination now wander from the benign chaos of Manhattan to the urgent logistical demands of a distant conflict, or the pressing envoy of messages to distant conversations that are aimed at preventing conflict. Surely, enlarging the number of supply channels can only improve the efficiency of supply to the front. But if the channels are congested it need not be so!

How relevant are the models?

An argument against the relevance to real situations of the four examples given here is that they are so simple. Examples closer to the complexity of real life, it might be argued, would contain fewer surprises. The simple models omit all the familiar features of real situations. According to this criticism, the surprises in these examples arise from the many respects in which the examples differ from reality rather than from those aspects the examples share with reality.

The force of this criticism depends on how well the models mirror key elements of complex real situations. If the criticism is intended in a constructive rather than dismissive mode, a constructive response is: propose better models, analyze them, and see how well they confirm intuitive rules of thumb.

In defending the study of the iterated prisoners' dilemma, Axelrod wrote: "It is the very complexity of reality which makes the analysis of an abstract interaction so helpful as an aid to understanding" (1984, p. 19). That argument applies here as well. Interactions in conflict and cooperation are too complex to analyze without the help of clear models.

These simple examples should at least shake any presumption that intuition alone suffices to anticipate the surprises of conflict and cooperation. The examples illustrate that the structure of interactive situations may lead to results that are surprises for all participants.

Joel E. Cohen has been professor of populations and head of the Laboratory of Population at the Rockefeller University since 1975. He received doctorates in applied mathematics (1970) and population sciences and tropical public health (1973) from Harvard University. His research deals mainly with the demography, ecology, and epidemiology of human and nonhuman populations and with mathematical concepts useful in these fields. His research is supported in part by the NSF. Address: Rockefeller University, New York, NY 10021–6399.

References

Arrow, K. J. 1963. *Social Choice and Individual Values*, 2nd ed. Wiley.

Axelrod, R. 1979. The rational timing of surprise. *World Politics* 31:228–46.

———. 1984. *The Evolution of Cooperation*. Basic Books.

Balinski, M. L., and H. P. Young. 1982. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale Univ. Press.

Bollobás, B. 1985. *Random Graphs*. Academic Press.

Bracken, P. 1983. *The Command and Control of Nuclear Forces*. Yale Univ. Press.

———. 1988. Do we really want to eliminate the chance of accidental war? *Defense Anal*. 4(1):81–89.

Braess, D. 1968. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* 12:258–68.

Brams, S. J., and P. C. Fishburn. 1983. *Approval Voting*. Birkhäuser.

Bunch, B. H. 1982. *Mathematical Fallacies and Paradoxes*. Van Nostrand Reinhold.

Carter, A. B., J. D. Steinbruner, and C. A. Zraket. 1987. *Managing Nuclear Operations*. Brookings.

Chung, F. R. K., J. E. Cohen, and R. L. Graham. 1987. Pursuit-evasion games on graphs. *J. Graph Theory* 12(2):159–67.

Cohen, J. E. 1986. Connectivity of finite anisotropic random graphs and directed graphs. *Math. Proc. Cambridge Philos. Soc*. 99:315–30.

Cohen, J. E., J. Hajnal, and C. M. Newman. 1986. Approaching consensus can be delicate when positions harden. *Stochastic Proc. and Appl*. 22:315–22.

Cohen, J. E., and F. P. Kelly. Unpubl. A paradox of congestion in a queuing network.

DeGroot, M. H. 1974. Reaching a consensus. *J. Am. Statist. Assoc*. 69: 118–21.

Downs, G. W., D. M. Rocke, and R. M. Siverson. 1985. Arms races and cooperation. *World Politics* 38:118–46. Repr. *Cooperation Under Anarchy*, ed. Henry S. Bienen, Richard D. Challener, Kenneth A. Oye, and Lynn T. White III, Princeton Univ. Press, 1986.

Dresher, M. 1961. *Games of Strategy: Theory and Applications*. Prentice-Hall.

French, J. R. P., Jr. 1956. A formal theory of social power. *Psychol. Rev*. 63:181–94.

Genest, C., and J. V. Zidek. 1986. Combining probability distributions: A critique and an annotated bibliography. *Statist. Sci*. 1(1):114–48.

Harary, F. 1959. A criterion for unanimity in French's theory of social power. In *Studies in Social Power*, ed. D. Cartwright, pp. 168–82. Ann Arbor: Inst. Soc. Res.

———. 1969. *Graph Theory*. Addison-Wesley.

Isaacson, D. L., and R. W. Madsen. 1976. *Markov Chains: Theory and Applications*. Wiley.

Kettelle, J. D. 1981. Conflict termination. *Proc. 48th Military Oper. Res. Soc*., pp. 7–17.

Machiavelli, N. 1521. *The Art of War*, trans. Ellis Farneworth, ed. Neal Wood. Bobbs-Merrill, 1965. See esp. book 7.

Maxwell, E. A. 1959. *Fallacies in Mathematics*. Cambridge Univ. Press.

Palmer, E. M. 1985. *Graphical Evolution: An Introduction to the Theory of Random Graphs*. Wiley.

Roberts, F. S. 1976. *Discrete Mathematical Models*. Prentice-Hall.

Robinson, D. R., and L. R. Foulds. 1980. *Digraphs: Theory and Technique*. Gordon & Breach.

Saari, D. G. 1987a. The source of some paradoxes from social choice and probability. *J. Econ. Theory* 41:1–22.

———. 1987b. A dictionary for voting paradoxes. Center for Mathematical Studies in Economics and Management Science, Discussion Paper 745. Northwestern Univ.

Steinberg, R., and W. I. Zangwill. 1983. The prevalence of Braess' paradox. *Transport. Sci*. 17(3):301–18.

Steinbruner, J. 1984. Launch under attack. *Sci. Am*. 250(1):37–47.

Szekély, G. J. 1987. *Paradoxes in Probability Theory and Mathematical Statistics*. Reidel.

Tuchman, B. 1962. *The Guns of August*. Macmillan.

von Neumann, J. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Ann*. 100:295–320. (On the theory of games of strategy, trans. Sonya Bargmann, in *Contributions to the Theory of Games*, ed. A. W. Tucker and R. D. Luce, vol. 4., pp. 13–42, Princeton Univ. Press, 1959.)

Zraket, C. 1984. Strategic command, control, communications, and intelligence. *Science* 224:1306–11.