The University of Pennsylvania Orthopaedic Journal
14: 75–77, 2001

Topics in Medical Economics: Lessons of the Prisoner's Dilemma*,†

JOSEPH BERNSTEIN, M.D., M.S.

The jousting between surgeons and insurance companies over fees can be fierce, and to the outside observer it seems anything but playful. Nevertheless, the insights that surgeons need to prevail in this battle may lie in a branch of theoretical economics called game theory.

The word *game* typically refers to a sport or contest, with the implicit connotation of recreation. The technical definition of a game, however, omits any notion of amusement. A game is defined as an encounter in which players execute plans of action under a set of rules to maximize their score—points, money, territory, whatever is at stake. Game theory, then, is the study of strategy and tactics. Game theory can be used to plan the moves of a chess game but can also be applied to any real-world situation in which the interaction between "players" satisfies the definition of a game. In fact, one of the first practical applications of game theory was far removed from the realm of fun and play: the Allied forces used game-theory methods to allocate resources and choose targets in World War II.

The negotiations between buyers and sellers (such as insurance companies and doctors, food-makers and farmers, or automobile manufacturers and steel mills), the posturing of military opponents, or even the debate between a couple over the evening's entertainment plan can take on the trappings of a game. Each side can make moves to maximize its "score."

Some scenarios in life and the strategies that they evoke can be reduced, in broad terms at least, to familiar game-theory concepts. When a situation under study resembles a particular game type or position, we say that the situation is isomorphic with the game. The usefulness of isomorphism is that, once a particular situation is recognized as a familiar game, many of the insights derived from careful study of the game can be applied to the situation at hand without repeated detailed analysis.

From my perspective, the current imbroglio between managed-care companies and surgeons over payment for surgical work resembles a classic game-theory problem called the prisoner's dilemma [3]. It should take no leap of the imagination to guess that it is the beleaguered doctor who plays the role of the prisoner here, yet I do not offer this comparison to make doctors feel even more besieged. Rather, I intend to share some of the analyses to resolve the prisoner's dilemma offered over the years to help doctors cope with their obviously difficult situation.

The prisoner's dilemma is the prototype of a class of games in which use of a reasoned and logical strategy (the so-called dominant solution) leads to a suboptimal result. (In their book *Decisions, Decisions: Game Theory and You,* Bell and Coplans [2] entitled the chapter on the prisoner's dilemma "Morons Do Better Than Logicians," and they may be right.) I believe that the use of logical reasoning by doctors responding to the advent of managed care has made their position worse. I further believe that this logic was not faulty: rather, it was the use of logic itself that caused the problem. Let's look at the prisoner's dilemma to see why.

The classic prisoner's dilemma goes as follows: Smith and Jones are both accused of grand larceny and are taken to jail and placed in separate cells. The district attorney approaches them individually and offers a deal. "We don't have enough evidence to convict you of grand larceny—only breaking and entering," she says, "so I offer you the following deal: turn state's evidence on your partner. If you confess, and he remains silent, I will let you go. But I warn you, if you remain silent, and he confesses, you will get the maximum ten-year sentence. Now, if you both confess, I can't let both of you go, so you each will get a five-year term, and if you both remain silent, I am sure we can lock you up for a year on the breaking-and-entering charge. By the way, I have offered the same deal to your partner and told him that I was talking to you."

The various combinations of responses and the sentences that the prisoners could receive are described in the table below, a so-called payoff matrix. (By tradition, acting in the interest of the other player is termed "cooperation" and acting against that interest is "defection," though one is not formally cooperating or defecting, since no communication is allowed (Table 1).)

The dilemma for each is stark: if only the prisoners could communicate—or, more to the point, trust the good intentions of the other—they would guarantee themselves a one-year sentence by remaining silent. But since they cannot communicate (and may not trust the other even if they could) they must employ logic, and logic gets them in trouble.

*No benefits in any form have been received or will be received from a commercial party related directly or indirectly to the subject of this article. No funds were received in support of this article.

†Reprinted with permission from *Journal of Bone and Joint Surgery* 2000; 82A:595–598.

From the Department of Orthopaedic Surgery, Leonard Davis Institute of Health Economics, 424 Stemmler Hall, University of Pennsylvania School of Medicine, Philadelphia, PA.

	Jones Confessses ("Defection")	Jones Remains Silent ("Cooperation")
Smith Confesses ("Defection")	Smith gets five years Jones gets five years	Smith goes free Jones gets ten years
Smith Remains Silent ("Cooperation")	Smith gets ten years Jones goes free	Smith gets one year Jones gets one year

Smith could say to himself, "What should I do? Well, let's see: maybe Jones will confess. In that case, I could confess too and get five years, or I could remain silent and get ten. If Jones confesses, confessing is clearly my best option. On the other hand, maybe Jones will remain silent. In that case, I could remain silent and we'd both get one year, or I could confess and go free. If Jones remains silent, confessing is the best option then too. Thus, no matter what Jones does, I am better off confessing."

This logic can be applied with equal validity by Jones, and he too will come to the same conclusion: confessing is better. Thus, using unassailable logic, both Smith and Jones will confess, and both will get five years. If only both had remained silent, they would have gotten only a one-year term, but logic would not let them. And, indeed, if Smith had unilaterally decided to remain silent, odds are that Jones would have confessed and Smith would be facing ten years instead of five. The logic is frustrating, but it is not wrong.

The prisoner's dilemma looks a lot like the situation offered to doctors by managed-care companies entering a new market. Consider, for example, a town with two orthopaedic surgeons, Dr. Smith and Dr. Jones. (This is the simple case, but the results can be generalized to a more realistic setting in which many doctors are at work.) Let's say that Dr. Smith and Dr. Jones each have 50% of the market and get paid the same fees for their work. They both earn a nice salary but would like to earn more, and each feels that he could care for additional patients. In fact, both think that they could double their effort.

A managed-care company can approach Dr. Smith and say, "Sign up with us. The rates we offer are lower than what you are used to, but if you sign up with us and Dr. Jones doesn't, you can have the entire market. You will make a lot more money. But take note: Dr. Jones may sign up. In that case, if you don't sign up with us, you will have nothing. Of course, if both of you sign up, your market share remains unchanged and your pay goes down. If both of you refuse to sign up, well, then we would be forced to pay you your usual fees, more than we currently offer."

Clearly, the doctors will be best off if they could agree to not sign, keeping things the way they are. But federal antitrust laws put Smith and Jones in separate cells—they are forbidden by law to collude—and they may not trust each other, to boot. Thus, the same logic that drives the prisoners to confess thrives here. Logic coerces both Dr. Smith and Dr. Jones to sign, with a net result of lower fees and no increase in market share.

Here is the payoff matrix (Table 2):

	Jones Signs with HMO ("Defection")	Jones Refuses to Sign ("Cooperation")
Smith Signs with HMO ("Defection")	Smith retains his original 50 percent of the market but at lower fees, and **Smith's income goes down** Jones retains his original 50 percent of the market but at lower fees, and **Jones's income goes down**	Smith gets 100 percent of the market. The lower fees are more than offset by higher volume, and **Smith's income goes up** Jones loses his patient base, and **Jones's income goes down substantially**
Smith Refuses to Sign ("Cooperation")	Smith loses his patient base, and **Smith's income goes down substantially** Jones gets 100 percent of the market. The lower fees are more than offset by higher volume, and **Jones's income goes up**	Smith retains his original 50 percent of the market at his usual rates, and **Smith's income remains the same** Jones retains his original 50 percent of the market at his usual rates, and **Jones's income remains the same**

Let's see what Dr. Smith would say to himself. 'What should I do? Maybe Dr. Jones will sign. In that case, I have two choices: to sign or to refuse. If I sign, at least I get to keep my patients, even if my income will go down. But if he signs and I don't, I get shut out. Clearly, if Dr. Jones will sign, I should sign too. What happens if Dr. Jones refuses to sign? In that case, I could refuse too, and we'll both do OK, keeping our market share with no loss of income. But I could sign and steal his market share. If I sign and he refuses, I'd make a lot more money. Signing, then, is the best option, independent of Dr. Jones's action."

This logic can be applied with equal validity by Dr. Jones, and he too will come to the same conclusion. Thus, using unassailable logic, both Dr. Smith and Dr. Jones will sign with the HMO and both will get lower fees without an increase in market share. If only both had refused to sign they would have gotten a much better deal, but logic would not let them. Thus, managed care offers doctors a prisoner's dilemma.

The game of the prisoner's dilemma can be made a little more interesting if it is played not once but many consecutive times by two players who can remember how the other player behaved in previous encounters. Here, a history of prior behavior and the opportunity for payback in the future may influence the choices that each player makes. Logic may insist that defecting (confessing to the district attorney or signing with the HMO) is the right thing to do if the choice is made only once, but if good behavior may be rewarded, or lack of cooperation may be punished, the correct choice may be different. This version of the game is called the iterated prisoner's dilemma, and it not only is subtler from the game-theory point of view, but it more closely resembles the situation that doctors encounter in real life. For example, contracts have to be renewed.

The iterated prisoner's dilemma invites a higher level of analysis by each player. One must not only calculate the effect that a choice has on the current payoff but also estimate the behavior that a given action will engender from the other player in future encounters. This new analysis was provided by Axelrod and Hamilton, [1] who reported their work in a landmark paper entitled "The Evolution of Cooperation" in the journal *Science* in 1981. The authors opened a contest to various strategies for the iterated prisoner's dilemma. Contestants sent in their strategy as a simple computer program, and all strategies encountered each other for an arbitrary number of rounds. Points were assigned for each interaction, corresponding to the length of the prison sentence in the original description. (The actual point values are arbitrary and can be positive or negative, as long as they are in the same relation as in the original.) The program with the highest point total was declared the winner.

To the surprise of many, a simple strategy named Tit-for-Tat was the victor. Tit-for-Tat does not machinate over its decisions; it does not use complex stochastic models. It simply cooperates with the other player in its first encounter—keeping silent or, if you prefer, refusing to sign with the HMO—and then subsequently acts exactly as the opponent did in the previous encounter. Over the long run, this turned out to be better than a strategy of always cooperating, always defecting, or any variant in between.

Even more striking was the outcome from a second contest conducted after results from the first were announced. In this one Tit-for-Tat was the winner. New contestants, informed that Tit-for-Tat had won the last time and explicitly told that it would be reentered in the contest, still could not beat it. The game theoreticians were impressed. They wondered what made Tit-for-Tat so good.

Tit-for-Tat, along with other strategies that also did well in the contest, was dissected, and four features were seen to be essential to winning an iterated prisoner's dilemma. The first is "niceness," which means, simply, refraining from defecting first. Defection tends to breed ill will from the other players and produces a course of mutual destruction. Being nice proved helpful. On the other hand, Tit-for-Tat was better than an "always cooperate" approach, indicating that provocability—that is, the ability to respond when the other player is not nice—is essential too; there is no point in being a dupe for others. And since players do better when both are cooperating, as compared with when both are defecting, the willingness to return to cooperation once the other side does first (a feature termed "forgiveness") is likewise important. The final feature gleaned was "clarity"—it must be obvious to other players that you are nice, provacable, and forgiving for those traits to serve you well.

The relevance of the prisoner's dilemma for the orthopaedic surgeon is, to my mind, uncontested. Managed-care companies approach surgeons with an offer that cannot be refused, or so it seems. They frame the option "sign with us!" in such a way that no other action seems reasonable. Signing with HMOs, in the absence of legalized (and enforceable) cooperation, seems to be the only way to survive. But clearly, if all doctors sign, none can gain market share and all become losers, and if no doctor signs, none can lose. Only the tantalizing prospect of cannibalizing a fellow doctor's practice or the fear of being eaten oneself obscures that undeniable fact.

When a managed-care company approaches with what seems to be an opportunity to kill or be killed, consider the payoff matrix. If you recognize a prisoner's dilemma lurking in it, remember that the best response has been proven. You must be nice and forgiving, yet provocable and clear. In your first encounter, at least, "cooperate" with your fellow doctor and assume that he or she will as well. Of course, retaliate when struck, but also be quick to return to niceness.

And one more thing: limit your envy to your fellow doctor's golf handicap. Envy can ruin a cooperative environment. In the iterated prisoner's dilemma contest, Tit-for-Tat did not outscore every individual strategy in one-on-one encounters; in fact, it did not outscore any. Its victory was based on the fact that in all such encounters Tit-for-Tat was satisfied with a tie—mutual success. That apparent meekness was enough to keep it on top overall. The lesson here is that surgeons must be satisfied to do only as well as (but no better than) their fellow doctors. The urge to outscore an opponent rather than to do as well as a colleague invites a downward spiral of defection and shared disadvantage.

Most orthopaedic surgeons are currently working harder than ever, and making less for it. How did we get there? Weren't all of our actions logical? Of course they were, but we did not realize that we were mired in a prisoner's dilemma and that logic is not the answer. So let's try a new way: cooperation. If we cooperate with each other, we may not trouce our rivals, but we won't get trounced ourselves. This strategy is not only nice, it is wise.

Note: The author thanks G. B. Holt, M.B.A., for his comments and insights.

References

1. Axelrod R, Hamilton WD: The evolution of cooperation. *Science,* 211:1390–1396, 1981.
2. Bell R, Coplans J: *Decisions, Decisions: Game Theory and You.* New York: Norton; 1976.
3. Hofstadter DR: The prisoner's dilemma computer tournaments and the evolution of cooperation. In *Metamagical Themas: Questions for the Essence of Mind and Pattern.* New York: Basic Books; 1985. p 715–734.