

American Economic Association

Puzzles: Slot Machines, Zomepirac, Squash, and More

Author(s): Barry Nalebuff

Source: *The Journal of Economic Perspectives*, Vol. 4, No. 1 (Winter, 1990), pp. 179-187

Published by: American Economic Association

Stable URL: <http://www.jstor.org/stable/1942839>

Accessed: 26/07/2010 15:30

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=aea>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Economic Perspectives*.

<http://www.jstor.org>

Puzzles

Slot Machines, Zomepirac, Squash, and More

Barry Nalebuff

This issue of puzzles can only be described as eclectic. The topics range from medical statistics to TV game shows to Las Vegas slot machines to squash to running around in squares. As usual, most puzzle answers are at the back. However, the first and last puzzles are unsolved. I offer a conjecture for one—the right answer could prove quite valuable. The column ends with a few comments on previous column and reader mail.

We continue to award T-shirts for the best new puzzles and most innovative answers. The standards for winning a T-shirt are quite high, but not impossible—the new winners are Avinash Dixit, John Geanakoplos, Dwight Jaffee, and David Lane. Keep sending your answers, comments, favorite puzzles and T-shirt size to me directly: Barry Nalebuff, “Puzzles,” Yale School of Organization and Management, Box 1A, New Haven, CT 06520.

Puzzle 1: The Two Bandits—Casinos and the Las Vegas Slots

Dwight Jaffee (Princeton) relates the latest attempt from Las Vegas to stimulate play on the slot machines. The major discouragement to playing slot machines is that the odds are against you. To counter this perception, some casinos have begun to advertise the payback ratio for their machines—the fraction of each dollar bet returned in prize money. Going one step further, one casino guarantees that they have some machines which are set to a payback ratio greater than 1!¹ If you could find

¹This advertisement about average paybacks and favorable machines applies only to the slot machines in a particular room, all of which have the same stake. If there were slot machines with different stakes, then we would need to know that the “average” was weighted by stake. Otherwise, the casino could offer the highest payback ratios on the 25 cent machines while depressing the dollar machines.

■ *Barry Nalebuff is Professor of Economics and Management, Yale University, SOM, New Haven, Connecticut.*

those machines and play them, you would expect to make money. The trick, of course, is that they don't tell you which machines are which. When they advertise that the average payback is 90 percent and that some machines are set at 120 percent, that also means that some other machines must be set somewhere below 90 percent. In addition, there is no guarantee that the same machines are set the same way each day—today's favorable machines could be tomorrow's losers. How might you go about guessing which machines are which?

Puzzle 2: The Bayesian Butler Did It

Dr. Michael Kramer, an epidemiologist at McGill, offers a fascinating case study that shows the difficulty of applying Bayesian inference to a real world problem (see Kramer, 1986). With the help of David Lane, a medical statistician at Wisconsin, I've simplified the medical features so that economists can try their hand.

The case concerns a basically healthy 42-year-old woman who suddenly died the afternoon following a midday visit to her dentist to have a wisdom tooth extracted. The coroner's report showed there was no evidence of violence. The most likely cause seemed to be an allergic reaction (or more technically speaking, anaphylactic shock) caused by an adverse drug reaction. Prior to visiting the dentist, the patient took a penicillin tablet (because of a history of a heart murmur) and thus we have one known suspect. The dentist prescribed a pain medication, zomepirac, which she was to take if needed. Evidence from the local pharmacy shows that the patient did indeed have the prescription filled. However, no one knows whether or not she ever took this medication for nobody counted the number of pills remaining in the bottle. Thus we have a second suspect, although there seems to be some crucial information missing. The third suspect is the novocaine used during the extraction procedure.

Dr. Kramer examined all the different competing risks and concluded that if we knew she had taken the zomepirac, then we would place a 95 percent probability that zomepirac was responsible for her death. You should take this conclusion as a statement of fact about the relative likelihood of the competing risks. But the problem is that we don't know if she took the drug. Fortunately, there is one additional piece of relevant information. A study of other patients who had similar operations shows that 60 percent had sufficient post-operative pain that they took the pain medication prescribed.

So what is the chance that zomepirac was responsible for her death? Is the answer simply the chance she took it times the chance it killed her: $(0.60)(0.95) = 0.57$?

Puzzle 3: Squash That Problem!

A running question over the last two columns concerns the differences in scoring between American and English squash. American squash is played to fifteen with every rally scored as a point while English squash is played to 9, but only the server scores points—if the receiver wins a rally, he is rewarded with the serve.

The question for tournament designers is: which system provides the better player with the biggest chance of winning? Outside the squash courts, economists can think more broadly about designing contest incentive schemes where effort and skill count more than luck. The perception from the *New Scientist* article (quoted in the Summer 1989 column) was that English rules were more favorable to the better player. The reason is that the receiver has to win two rallies in a row to score a point and this is much harder for the weaker player. Is that correct? For simplicity, you may assume that American squash played to fifteen points takes the same time to complete a game as scoring with the British rules to nine.

Puzzle 4: Running Around in Squares

This problem calls for some cleverness. But before I describe the puzzle at hand, it will help put you in the right frame of mind if I first relate a well-known problem concerning two trains and a fly. The two trains are initially 100 miles apart and are approaching each other at 50 miles per hour. The fly flies back and forth at 100 mph between two oncoming trains. The question is how far does the fly fly before it is squashed by the two trains?

There are two ways to solve this problem. One is to sum up an infinite series. On the first leg, the fly starts at one train and flies $66\frac{2}{3}$ miles before running into the other train and turning around. But now the trains are only $33\frac{1}{3}$ miles apart so the fly's next leg is only $\frac{1}{3}$ the previous distance. Thus the total distance travelled is simply the sum of the infinite series $66.66[1 + \frac{1}{3} + \frac{1}{9} + \frac{1}{27} + \dots] = 66.66(\frac{3}{2}) = 100$. The cleverer way to solve the problem is to see that it takes the two trains exactly one hour to meet and a fly travelling at 100 mph covers 100 miles during that time.

The fame of this problem among economists comes from a story about John von Neumann. He was given this problem at a cocktail party and instantaneously responded with the answer: 100 miles. The person posing the question complimented von Neumann on recognizing the clever solution, for his experience showed that most people labor to solve the infinite series. Professor von Neumann responded: But that's how I did it!

After being told this story, Richard Palmer (a polymath at Duke) offered a physicist's version. There are four individuals, located at the four corners of a 100 by 100 square. Each person begins walking towards the person to their right. How far do they walk before meeting in the center?

The hard way to solve the puzzle is to integrate the distance along the spiral path they follow. Of course, there is an easier way. The only thing more impressive than figuring out the easy solution instantaneously is to figure it out the hard way just as quickly.

Avinash Dixit (Princeton) showed off his Cambridge Math Tripos training and solved it right away. Even better, he offered a variation of the puzzle which makes it

seem even harder, although those who figure out the quick solution will know how to solve this too. Instead of locating four people at the corners of a square, what if there are N people located at the vertices of a symmetric n -gon? To give you a hint, the answer for an octagon is $200/[2 - \sqrt{2}]$.

Puzzle 5: For 40 Points, This TV Game Show is the Most Challenging

What is “Jeopardy”? For the uninitiated, this TV game show has three contestants who compete to figure out the questions based on the presented answer. Each time a contestant provides another right question, her score increases by the question’s value. Deductions are made for incorrect questions.

The climax of the show is called Final Jeopardy. In this section, the three contestants are told the subject matter and then asked to prespecify a wager amount. They are allowed to bet any amount up to their total score at even odds. The person who ends up with the highest total gets to keep the score as cash and is invited back to defend against two new challengers the next day. The others are given small consolation prizes. Here lies an unsolved problem: how much should a contestant bet?

Sometimes the answer is easy. If the leader has more than double the next highest person, then a small bet guarantees a win. Unless the person is super-confident about the subject matter, the option value of the returning the next day makes the guaranteed win an extremely attractive option.²

The problem is what to do when the scores are closer. If the two front-runners both bet large amounts and get the question wrong that their final scores may end up quite low. In this case, the contestant with the lowest score may emerge victorious simply by betting a small amount.

It’s safe to say that the optimal betting strategy depends critically on the contestant’s belief that she can question the answer and her prior belief about how likely the others are to know the right question. For simplicity, let us assume that these probabilities are equal to their observed success rate so far, and thus common knowledge. Next, we know that the betting strategies will depend on the scores going into final jeopardy. The only element remaining is to specify the objective functions. Here too we simplify to say that all the contestants care about is winning (concern over the size of their winnings makes the problem too complicated). Now and then I watch Jeopardy to see what the contestants do. People who are amazingly clever with information about Fidel Castro’s brother or the geography of Africa don’t fare as well when it comes to gambling in Final Jeopardy. What advice can economists offer?

²Things get more complicated if a contestant wins five days in row. In that case, they are retired as an undefeated champion. It may be worth taking the wager given a high enough prior belief in being able to answer the question.

Answers to Puzzles

Answer to Puzzle 1

Sadly, Dwight didn't have an answer and neither do I. If we did, I can assure you we could have used it to test the law of large numbers. Nonetheless, I'm willing to offer a conjecture. The trick is to put yourself into the casino owner's shoes. They only make money provided that people play the disadvantageous machines at least as much as the favorable machines. Hence, you can bet that the machines most heavily played are not the ones with the highest payback.

But wait! If the payback is what attracts people to the machines in the first place, how can they hide the machines that are offering the best odds? If people play the machines which pay out the most, won't they find the best ones? Not necessarily and especially not necessarily in time!

The payoff of the machine is in large part determined by the chance of a jackpot prize. Look at a slot machine that takes a quarter a pull. A jackpot prize of \$10,000 with chance 1 in 40,000 would give a payoff ratio of 1. If the casino raised the chance to 1 in 30,000 then the payoff ratio would be very favorable, at 1 to 33. But people watching others play the machines would almost always see a person dropping quarter after quarter with no success. A very natural conclusion would be that this is one of the least favorable machines. Eventually, when the machine pays its jackpot prize, it could be retooled and set at a lower rate.

In contrast, the least favorable machines could be set to pay back a large fraction of the money with a very high chance and basically eliminate the hope of the big jackpot. Look at a machine set with a payback of 80 percent. If it provided a \$1.00 prize on roughly every fifth draw then this machine would make a lot of noise, attract attention, and possibly more gamblers' money.

Perhaps the experienced slot players have figured all this out. If so, you can bet that the casinos are just doing the reverse. Whatever happens, the casinos can find out at the end of the day which machines were played the most. They can make sure that the payoff patterns that attract the most play are actually the ones with the lower payoff ratio. For while the difference between a payoff ratio of 1.20 and 0.80 may seem large—and determines the difference between making money and losing money—it can be extremely hard to distinguish based on a number of experiments any one slot player has to make. The casinos can design the payoffs to make these inferences harder and even go the wrong way most of the time. Your best bet is to recognize that unlike the United Way, the Las Vegas casinos are not in the business to give out money. In their search for favorable machines, the majority of the players can't be right. If the majority of the people were able to figure this all out, then the casino would discontinue their offer. Enjoy playing the slot machines for their entertainment value, for that's the surest payoff.

Carl Shapiro adds a New Jersey perspective with some interesting facts about slot machines in Atlantic City. First, it is illegal to advertise payback ratios in New Jersey. (Are the regulators worried that competition would lead to higher payback ratios and

more slot machine play? Would this be worse for the gamblers or for the casinos?) In Atlantic City, the slots near the doors and other prominent places have the highest payback ratios. Perhaps they also are tuned differently so as to pay back more frequently (and more noisily) but with smaller prizes.

Answer to Puzzle 2

The answer .57 is awfully tempting but completely wrong. The reason is that it ignores the fact that this patient, unlike most other patients, died. This information should be used to update our prior belief about the likelihood she took the prescribed pain-killer. On the other hand, to apply the standard Bayesian rule, we need to know the unconditional probability that the patient died and that number doesn't seem immediately available. Here is where some creativity is needed. Consider a group of 1000 patients. Of the 600 who take the pain medication, say that X people die. How many will die from the group of 400 who didn't take the pain medication? For the first group, we are told that $0.95X$ died from the drug and $0.05X$ died from other causes. The people who died from other causes should be just as likely to die if they didn't take the drug as if they did! For example, the chance of a fatal bee sting should not depend on whether or not the pain medication was used. (Actually, the pain medication could have some synergistic effects that either increase or decrease the other competing risks, but it seems reasonable to ignore this possibility.) Thus, the best guess for the number of deaths in the group of 400 who didn't take the pain medication is $400(0.05X/600)$. Now all the necessary information is on the table. The number of deaths due to zomepirac compared to all deaths is $0.95X/[X + (2/3)0.05X] = 95/103.3$, which is approximately 92 percent.

The odds are very strong that she took the drug and that the drug killed her. For those of you who are wondering whether this seems like a pretty dangerous medication, take comfort in the fact zomepirac has been taken off the U.S. market.

Answer to Puzzle 3

Actually, given that the two scoring systems lead to games that involve roughly the same number of rallies, there should be no difference between the two in favoring the better player. The reason is that all the unscored points in English squash cancel out. The easiest way to see this is to pretend that every rally in English squash was scored. When the receiver wins a rally, that's one point she wouldn't otherwise get. But then when she loses a rally when serving, it's one point she wouldn't otherwise lose. Every time a player loses a serve and gets it back is simply one point added to both scores. When determining who wins, this effectively cancels out. There may be some slight correction for the fact that there may be a temporary imbalance giving the initial server a one-point edge. But the important point is that whoever is ahead at the end is the one who wins and since the missing points essentially cancel each other out, the scoring rules don't have any significant influence on who is leading.

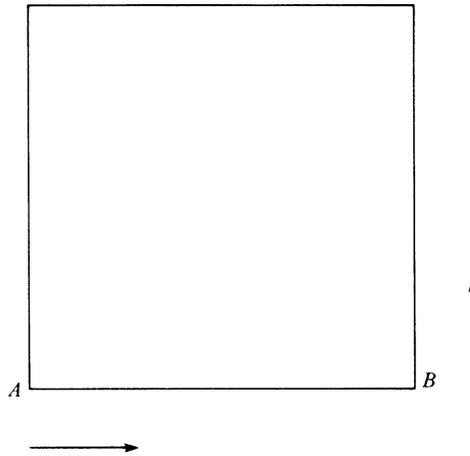
Thus the problem reduces to which game is longer. A longer game gives the better player a greater chance of being ahead. Here is where the simplifying assumption that games played under the different scoring rules take the same time

comes into play. More generally, one could adjust the winning levels to make this true or for the current levels simply find out which one takes longer to play and recognize that that's the one favoring the better player.

If instead it were always necessary to win two rallies in a row in order to score a point, then the unscored points would no longer cancel out. The better player would have a great advantage. When the better player's chance of winning a rally is $p > 1/2$, the odds of winning two rallies in a row (and thus a point) are p^2 : $(1 - p)^2$. For example, a player with a 20 percent edge ($p = 0.60$) finds the advantage roughly doubled—the player's chance of scoring a point grows to 69 percent and the chance of winning the game grows even more. The intuition that a rule which requires winning two rallies in a row before scoring helps the better player is correct provided that is true for all rallies, and not just the transition from server to receiver.

Answer to Puzzle 4

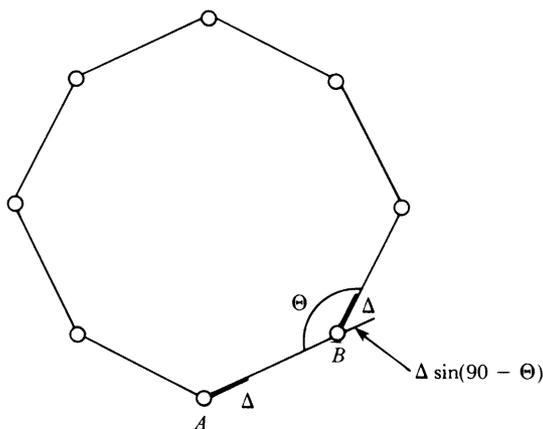
The trick is to look at the problem from the frame of reference of one of the individuals, say person *A*. From *A*'s perspective, the person to her right, *B*, is always walking at a right angle to the line connecting *A* and *B*. Thus, to the first order, *B*'s movement has no effect on the distance between *A* and *B*. Since *A* is continuously adjusting her direction towards *B*, it is always as if *A* is walking directly towards *B*.



Initially *A* and *B* are a distance of 100 apart. Thus *A* walks 100 before meeting *B* at the center.

More generally, the picture below shows that the distance travelled when there are n participants beginning at the vertices of a symmetric n -gon equals the side length over $[1 + \sin(90 - \theta)]$, where $\theta = (n - 2)180/n$, the measure of an interior angle for a symmetric n -gon. As *A* moves an infinitesimal amount Δ towards her right-hand neighbor *B*, *B*'s Δ movement towards her right-hand neighbor *C* contributes an amount $\Delta \sin(90 - \theta)$ towards the distance between *A* and *B*. Thus the

distance A must travel is scaled by the factor $1/[1 + \sin(90 - \theta)]$. Note that the scaling factor is smallest for a triangle ($2/3$), is unity for a square, and rises to infinity as the number of sides, n , approaches infinity. The reason is straightforward. In the limiting case, it is as if everyone is positioned on an infinitely large circle. It appears that you and your neighbor are walking in exactly the same direction and hence you never get any closer: an infinitely large circle has the same property as an infinite line!



Answer to Puzzle 5

I don't know.

Reader Mail and Comments on Previous Puzzles

In the giving credit where credit's due department, the solution I described to Rafael Hassin's waiting line problem is due to John Geanakoplos (Yale). Hassin showed by construction that a LIFO allocation rule resulted in efficient line lengths. As he presented this result at an IMSSS workshop, Geanakoplos offered an intuitive argument generalizing the result: any rule where the most recent arrival is placed at any position other than at the end of the line leads to efficient line lengths since the last person in line (and marginal waiter) imposes no externality on anyone who comes later. Geanakoplos's argument is the one presented in the Summer 1989 column.

In the better-late-than-never department, Dr. Aris Ikkos (Alexander, Young & Co. in Greece) provided a correct solution to the "Keeping One Step Ahead" puzzle published in Winter 1988. John Baffes (Univ. of Maryland, College Park) offered one more perspective on Joe Farrell's puzzle about competitive equilibrium as a welfare minimum. His point was that if the government subsidizes some goods and taxes others, the welfare optimization problem looks convex—unless appropriate constraints

are added. In a closed economy, supply must equal demand as otherwise there will be rationing. In an open economy, supply need not equal demand, but there must be a balanced budget. Adding this constraint restores the concavity of the maximization problem and thus the optimality of the competitive equilibrium.

■ *Once again my thanks go to Carl Shapiro and Tim Taylor for keeping me honest; they test and improve each puzzle before you see it.*

References

Kramer, Michael, "A Bayesian Approach to Assessment of Adverse Drug Reactions: Evaluation of a Case of Fatal Anaphylaxis," *Drug Information Journal*, 1986, 20, 505–521.