Games People Don't Play

Peter Winkler, Bell Labs, Murray Hill NJ 07974

Abstract

Not all games are to play; some of the most amusing are designed just to think about. Is the game fair? What's the best strategy? The games we describe below were collected from various sources by word of mouth, but thanks to readers of an earlier version we have written sources for some of them. An odd (actually, even) feature of the games in this article is that each has two versions, with entertaining contrasts between the two. There are four pairs of games: the first involving numbers, the second hats, the third cards, and the fourth gladiators. We present all the games first, then their solutions.

1 The Games

1.1 Larger or Smaller (standard version)

We begin with a classic game which makes a great example in a class on randomized algorithms. Paula (the perpetrator) takes two slips of paper and writes an integer on each. There are no restrictions on the two numbers except that they must be different. She then conceals one slip in each hand.

Victor (the victim) chooses one of Paula's hands, which Paula then opens, allowing Victor to see the number on that slip. Victor must now guess whether that number is the larger or the smaller of Paula's two numbers; if he guesses right he wins \$1, otherwise he loses \$1.

Clearly Victor can achieve equity in this game, merely by flipping a coin to decide whether to guess "larger" or "smaller". The question is: not knowing anything about Paula's psychology, is there any way he can do better than break even?

1.2 Larger or Smaller (random version)

Now let's make things much easier for Victor: instead of being chosen by Paula, the numbers are chosen independently at random from the uniform distribution on [0,1] (two outputs from a standard random number generator will do fine).

To compensate Paula, we allow her to examine the two random numbers and *to decide which one Victor will see.* Again, Victor must decide whether the number he sees is the larger or smaller of the two, with \$1 at stake. Can he do better than break even? What are his and Paula's best (i.e. "equilibrium") strategies?

1.3 Colored Hats (simultaneous version)

Each member of a team of n players is to be fitted with a red or blue hat; each player will be able to see the colors of the hats of his teammates, but not the color of his own hat. No communication will be permitted. At a signal each player will simultaneously guess the color of his own hat; all the players who guess wrong are subsequently executed.

Knowing that the game will be played, the team has a chance to collaborate on a strategy (that is, a set of schemes—not necessarily the same for each player—telling each player which color to guess, based on what he sees). The object of their planning is to *guarantee* as many survivors as possible, assuming worst-case hat distribution.

In other words, we may assume the hat-distributing enemy knows the team's strategy and will do his best to foil it. How many players can be saved?

1.4 Colored Hats (sequential version)

Again, each of a team of n players will be fitted with a red or blue hat; but this time the players are to be arranged in a line, so that each player can see only the colors of the hats in front of him. Again each player must guess the color of his own hat, and is executed if he is wrong; but this time the guesses are made sequentially, from the back of the line toward the front. Thus, for example, the ith player in line sees the hat-colors of players $1, 2, \dots, i-1$ and hears the guesses of players $i+1, \dots, n$ (but he isn't told which of those guesses were correct—the executions take place later).

As before, the team has a chance to collaborate beforehand on a strategy, with the object of guaranteeing as many survivors as possible. How many players can be saved in the worst case?

1.5 Next Card Red

Paula shuffles a deck of cards thoroughly, then plays cards face up one at a time, from the top of the deck. At any time Victor can interrupt Paula and bet $1 that the next card will be red. (If he never interrupts, he's automatically betting on the last card.)

What's Victor's best strategy? How much better than even can he do? (Assume there are 26 red and 26 black cards in the deck.)

1.6 Next Card Color Betting

Again Paula shuffles a deck thoroughly and plays cards face up one at a time. Victor begins with a bankroll of \$1, and can bet any fraction of his current worth, prior to each revelation, on the color of the next card. He gets even odds regardless of the current composition of the deck. Thus, for example, he can decline to bet until the last card, whose color he of course knows, then bet everything and be assured of going home with \$2.

Is there any way Victor can *guarantee* to finish with more than \$2? If so, what's the maximum amount he can assure himself of winning?

1.7 Gladiators, with Confidence-Building

Paula and Victor each manage a team of gladiators. Paula's gladiators have strengths p_1, $p_2 \ldots p_m$ and Victor's, $v_1 \ldots v_n$. Gladiators fight one-on-one to the death, and when a gladiator of strength x meets a gladiator of strength y, the former wins with probability $x/(x+y)$ and the latter with probability $y/(x+y)$. Moreover, if the gladiator of strength x wins he gains in confidence and inherits his opponent's strength, so that his own strength improves to $x+y$; similarly, if the other gladiator wins, his strength improves from y to $x+y$.

After each match, Paula puts forward a gladiator (from those on her team who are still alive), and Victor must choose one of his to face Paula's. The winning team is the one which remains with at least one live player.

What's Victor's best strategy? In particular, if Paula begins with her best gladiator, should Victor respond from strength or weakness?

1.8 Gladiators, with Constant Strength

Again Paula and Victor must face off in the Coliseum, but this time confidence is not a factor and when a gladiator wins he keeps the same strength he had before.

As before, prior to each match, Paula chooses her entry first. What is Victor's best strategy? Whom should he play if Paula opens with her best man?

2 Solutions and Comments

2.1 Larger or Smaller (standard version)

As far as we know, this problem originated with Tom Cover in 1986 and appears as a 1-page "chapter" in his book [2]. Amazingly, there *is* a strategy which guarantees Victor a better than 50% chance to win.

Before playing, Victor selects a probability distribution on the integers which assigns positive probability to each integer. (For example, he plans to flip a coin until a "head" appears. If he sees an even number $2k$ of tails, he will select the integer k; if he sees $2k-1$ tails, he will select the integer $-k$.)

If Victor is smart he will conceal this distribution from Paula, but as you will see Victor gets his guarantee even if Paula finds out.

After Paula picks her numbers, Victor selects an integer from his probability distribution and adds $1/2$ to it; that becomes his "threshold" t. For example, using the distribution above, if he flips 5 tails before his first head, his random integer will be -3 and his threshold t will be $-2\frac{1}{2}$.

When Paula offers her two hands, Victor flips a *fair* coin to decide which hand to choose, then looks at the number in that hand. If it exceeds t, he guesses that it is the larger of Paula's numbers; if it is smaller than t, he guesses that it is the smaller of Paula's numbers.

So why does this work? Well, suppose that t turns out to be larger than either of Paula's numbers; then Victor will guess "smaller" regardless of which number he gets, and thus will be right with probability exactly $1/2$. If t undercuts both of Paula's numbers, Victor will inevitably guess "larger" and will again be right with probability $1/2$.

But, *with positive probability*, Victor's threshold t will fall *between* Paula's two numbers; and then Victor wins regardless of which hand he picks. This possibility, then, gives Victor the edge which enables him to beat 50%.

Comment: Neither this nor any other strategy enables Victor to guarantee, for some fixed ε, a probability of winning greater than $50\% + \varepsilon$. A smart Paula can choose randomly two consecutive multi-digit integers, and thereby reduce Victor's edge to a smidgeon.

2.2 Larger or Smaller (random version)

It looks like the ability to choose which number Victor sees is paltry compensation to Paula for not getting to pick the numbers, but in fact *this* version of the game is strictly fair: Paula can prevent Victor from getting any advantage at all.

Her strategy is simple: look at the two random real numbers, then feed Victor the one which is closer to $1/2$.

To see that this reduces Victor to a pure guess, suppose that the number x revealed to him is between 0 and $1/2$. Then the unseen number is uniformly distributed in the set $[0,x] \cup [1-x,1]$ and is therefore equally likely to be smaller or greater than x. If $x > 1/2$ then the set is $[0,1-x] \cup [x,1]$ and the argument is the same.

Of course Victor can guarantee probability $1/2$ against any strategy by ignoring his number and flipping a coin, so the game is completely fair.

Comment: This amusing game was brought to my attention only a year ago, at a restaurant in Atlanta. Lots of smart people were stymied, so if you failed to spot this nice strategy of Paula's, you're in good company.

2.3 Colored Hats (simultaneous version)

It is not immediately obvious that any players can be saved. Often the first strategy considered is "guessing the majority color"; e.g. if $n = 10$, each player guesses the color he sees on 5 or more of his 9 teammates. But this results in 10 executions if the colors are distributed 5-and-5, and the most obvious modifications to this scheme also result in total carnage in the worst case.

However, it is easy to save $\lfloor n/2 \rfloor$ players by the following device. Have the players pair up (say, husband and wife); each husband chooses the color of his wife's hat, and each wife chooses the color she *doesn't* see on her husband's hat. Clearly, if a couple have the same color hats, the husband will survive; if different, the wife will survive.

To see that this is best possible, imagine that the colors are assigned uniformly at random (e.g. by fair coin-flips), instead of by an adversary. Regardless of strategy, the probability that any particular player survives is exactly $1/2$; therefore the expected number of survivors is exactly $n/2$. It follows that the *minimum* number of survivors cannot exceed $\lfloor n/2 \rfloor$.

2.4 Colored Hats (sequential version)

This version of the hats game was passed to me by Girija Narlikar of Bell Labs, who heard it at a party (the previous version was my own response to Girija's problem, but has no doubt been considered many times before). For the sequential version it is easy to see that $\lfloor n/2 \rfloor$ can be saved; for example, players n, $n-2$, $n-4$ etc. can each guess the color of the player immediately ahead, so that players $n-1$, $n-3$ etc. can echo the most recent guess and save themselves.

It seems like some probabilistic argument such as provided for the simultaneous version should also work here, to show that $\lfloor n/2 \rfloor$ is the most that can be saved. Not so: in fact, all the players except the last can be saved!

The last player (poor fellow) merely calls "red" if he sees an odd number of red hats in front of him, and "blue" otherwise. Player $n-1$ will now know the color of his own hat; for example, if he hears player n guess "red" and sees an *even* number of red hats ahead, he knows his own hat is red.

Similar reasoning applies to each player going up the line. Player i sums the number of red hats he sees and red guesses he hears; if the number is odd he guesses "red", if even he guesses "blue", and he's right (unless someone screwed up).

Of course the last player can never be saved, so $n-1$ is best possible.

2.5 Next Card Red

It looks as if Victor can gain a small advantage in this game by waiting for the first moment when the red cards in the remaining deck outnumber the black, then making his bet. Of course, this may never happen and if it doesn't, Victor will lose; does this compensate for the much greater likelihood of obtaining a small edge?

In fact it's a fair game. Not only has Victor no way to earn an advantage, he has no way to lose one either: all strategies are equally effective and equally harmless.

This fact is a consequence of the martingale stopping time theorem, and can also be established without much difficulty by induction (by two's) on the number of cards in the deck. But there is another proof, which I will describe below, and which must surely be in "the book"[1].

Suppose Victor has elected a strategy S, and let us apply S to a slightly modified version of "Next Card Red". In the new version, Victor interrupts Paula as before, but this time he is betting not on the *next* card in the deck, but instead on the *last* card of the deck.

Of course, in any given position the last card has precisely the same probability of being red as the next card. Thus the strategy S has the same expected value in the new game as it did before.

But, of course, the astute reader will already have observed that the new version of "Next Card Red" is a pretty uninteresting game; Victor wins if the last card is red, regardless of his

[1]As many readers will know, the late, great mathematician Paul Erdős often spoke of a book owned by God in which is written the best proof of each theorem. I imagine Erdős is reading the book now with great enjoyment, but the rest of us will have to wait.

strategy.

There is a discussion of "Next Card Red" in Tom Cover's book [3] on information theory, based on an unpublished result in [1].

Comment: The modified version of "Next Card Red" reminds me of a game which was described—for satiric purposes—in the *Harvard Lampoon*[2] many years ago. Called "The Great Game of Absolution and Redemption", it required that the players move via dice rolls around a Monopoly-like board, until everyone has landed on the square marked "DEATH". So how do you win?

Well, at the beginning of the game you were dealt a card face down from the Predestination Deck. At the conclusion you turn your card face up, and if it says "damned", you lose.

2.6 Next Card Color Betting

Finally, we have a really good game for Victor. But can he do better than doubling his money, regardless of how the cards are distributed?

It is useful first to consider which of Victor's strategies are optimal in the sense of "expectation". It is easy to see that as soon as the deck comes down to all cards of one color, Victor should bet everything at every turn for the rest of the game; we will dub any strategy which does this "reasonable." Clearly, every optimal strategy is reasonable.

Surprisingly, the converse is also true: no matter what Victor's strategy is, as long as he comes to his senses when the deck becomes monotone, his expectation is the same! To see this, consider first the following *pure* strategy: Victor imagines some fixed specific distribution of red and black in the deck, and bets *everything he has* on that distribution *at every turn.*

Of course, Victor will nearly always go broke with this strategy, but if he wins he can buy the earth—his take-home is then $2^52 \times \$1$, around 50 quadrillion dollars. Since there are $\binom{52}{26}$ ways the colors can be distributed in the deck, Victor's mathematical expected return is $\$2^{52}/\binom{52}{26} = \$9.0813$.

Of course, this strategy is not realistic but it is "reasonable" by our definition, and most importantly, *every reasonable strategy is a combination of pure strategies of this type.* To see this, imagine that Victor had $\binom{52}{26}$ people working for him, each playing a different one of the pure strategies.

[2] *Harvard Lampoon* Vol. CLVII No. 1, March 30, 1967, pp. 14–15. The issue is dubbed "Games People Play Number" and the particular game in question appears to have been composed by D.C. Kenney and D.C.K. McClelland.

We claim that every reasonable strategy of Victor's amounts to distributing his original \$1 stake among these assistants, in some way. If at any time his collective assistants bet $\$x$ on "red" and $\$y$ on black, that amounts to Victor himself betting $\$(x-y)$ on "red" (when $x > y$) or $\$(y-x)$ on black (when $y > x$).

Each reasonable strategy yields a distribution, as follows. Say Victor wants to bet \$.08 that the first card is red; this means that the assistants who are guessing "red" first get a total of \$.54 while the others get only \$.46. If, on winning, Victor plans next to bet \$.04 on black, he allots \$.04 more of the \$.54 total to the "red-black" assistants than to the "red-red" assistants. Proceeding in this manner, eventually each individual assistant has his assigned stake.

Now, any ("convex") combination of strategies with the same expectation shares that expectation, hence every reasonable strategy for Victor has the same expected return of \$9.08 (yielding an expected profit of \$8.08). In particular all reasonable strategies are optimal.

But one of these strategies *guarantees* \$9.08; namely, the one in which the \$1 stake is divided equally among the assistants. Since we can never guarantee more than the expected value, this is the best possible guarantee.

Comment: This strategy is actually quite easy to implement (assuming as we do that U.S. currency is infinitely divisible). If there are b black cards and r red cards remaining in the deck, where $b \geq r$, Victor bets a fraction $(b-r)/(b+r)$ of his current worth on black; if $r > b$, he bets $(r-b)/(r+b)$ of his worth on red.

If the original \$1 stake is *not* divisible, but is composed of 100 indivisible cents, things become more complicated and it turns out that Victor does about a dollar worse. A dynamic program (written by Ioana Dumitriu of M.I.T.) shows that optimal play by Victor and Paula results in Victor ending with \$8.08; Figure 1 shows the size of Victor's bankroll at each stage of a well-played game. For example, if the game reaches a point when there are 12 black and 10 red cards remaining, Victor should have \$1.08. By comparing the entries above and to the right we see that he should bet either \$.11 (in which case Paula will let him win) or \$.12 (in which case he will lose) that the next card is black.

Note that Victor tends to bet slightly more conservatively in the "100 cents" game than in the continuous version. If instead he chooses to bet always the nearest number of cents to the fraction $(b-r)/(b+r)$ of his current worth, Paula will knock him down to \$0 before half the deck is gone!

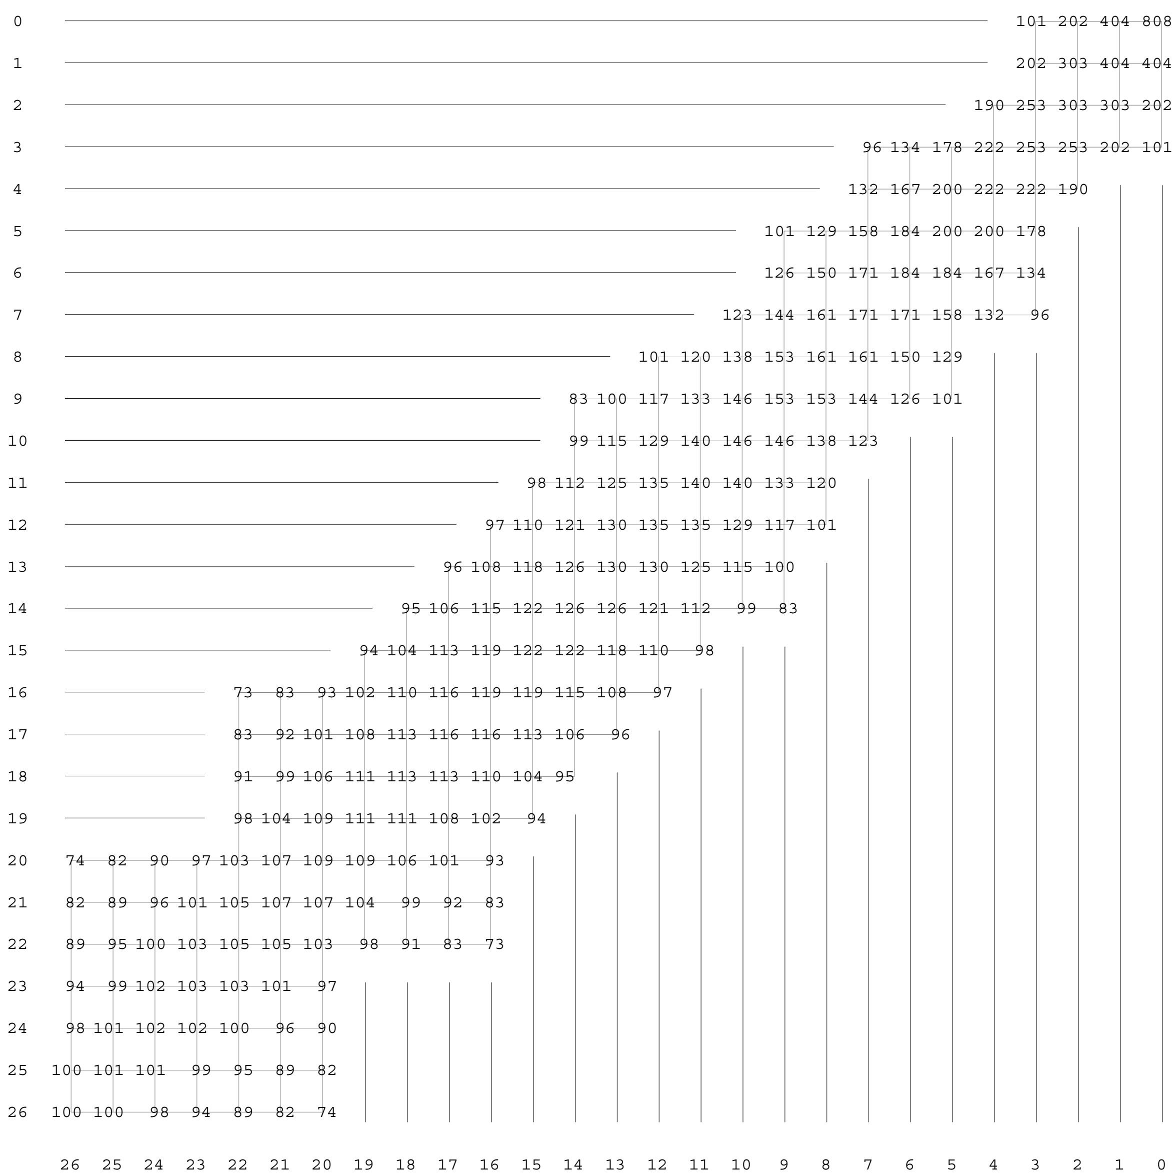

Figure 1: Optimal strategy in the discrete 100 cent Next Card Color game

I heard this problem from Russ Lyons, of Indiana University, who heard it from Yuval Peres, who heard it from Sergiu Hart; Sergiu doesn't remember where he heard it but suspects that Martin Gardner may have written about it decades ago!

2.7 Gladiators, with Confidence-Building

As in "Next Card Red", all strategies for Victor are equally good.

To see this, imagine that strength is money. Paula begins with $P = p_1 \dots p_m$ dollars and Victor with $V = v_1 \dots v_n$ dollars. When a gladiator of strength x beats a gladiator of strength y, the former's team gains $\$y$ while the latter's loses $\$y$; the total amount of money always remains the same. Eventually, either Paula will finish with $\$P + \$V$ and Victor with zero, or the other way 'round.

The key observation is that every match is a fair game. If Victor puts up a gladiator of strength x against one of strength y, then his expected financial gain is

$$\frac{x}{x+y} \cdot \$y \;+\; \frac{y}{x+y} \cdot (-\$x) \;=\; \$0\ .$$

Thus the whole tournament is a fair game, and it follows that Victor's expected worth at the conclusion is the same as his starting stake, $\$P$. Thus

$$q(\$P + \$V) + (1-q)(\$0) = \$P$$

where q is the probability that Victor wins. Thus $q = P/(P+V)$, independent of anyone's strategy in the tournament.

Comment: Here's another, more combinatorial, proof, pointed out by one of my favorite collaborators, Graham Brightwell of the London School of Economics. Using approximation by rationals and clearing of denominators, we may assume that all the strengths are integers. Each gladiator is assigned x balls if his initial strength is x, and all the balls are put into a uniformly random vertical order. When two gladiators battle the one with the higher topmost ball wins (this happens with the required $x/(x+y)$ probability) and the loser's balls accrue to the winner.

The surviving gladiator's new set of balls is again uniformly distributed in the vertical order, just as if he had started with the full set; hence the outcome of each match is independent of previous events, as required. But regardless of strategy, Victor will win if and only if the top ball in the whole order is one of his; this happens with probability $P/(P+V)$.

2.8 Gladiators, with Constant Strength

Obviously, the change in rules makes strategy considerations in this game completely different from the previous one—or does it? No, again the strategy makes no difference!

For this game we take away each gladiator's money (and balls), and turn him into a lightbulb.

The mathematician's ideal lightbulb has the following property: its burnout time is completely memoryless. That means that knowing how long the bulb has been burning tells us absolutely nothing about how long it will continue to burn.

The unique probability distribution with this property is the exponential; if the expected (average) lifetime of the bulb is x, then the probability that it is still burning at time t is $e^{-t/x}$.

Given two bulbs of expected lifetimes x and y, respectively, the probability that the first outlasts the second is—you guessed it—$x/(x+y)$. We imagine that the matching of two gladiators corresponds to turning on the corresponding lightbulbs until one (the loser) burns out, then turning off the winner until its next match; since the distribution is memoryless, the winner's strength in its next match is unchanged.

During the tournament Paula and Victor each have exactly one lightbulb lit at any given time; the winner is the one whose total lighting time (of all the bulbs/gladiators on her/his team) is the larger. Since this has nothing to do with the order in which the bulbs are lit, the probability that Victor wins is independent of strategy. (Note: that probability is a more complex function of the gladiator strengths than in the previous game).

Comment: The constant-strength game appears in [4]. I have a theory that the other game came about in the following way: someone enjoyed the problem and remembered the answer (all strategies equally good) but not the conditions. When he or she tried to reconstruct the rules of the game, it was natural to introduce the inherited-strength condition in order to make a martingale.

Afterword

It seems only right that I conclude with a game whose solution is left to the reader. Described by Todd Ebert in his 1998 U.C. Santa Barbara PhD thesis,[3] it can be thought of as another version of the (simultaneous) colored hats game. In this version the colors really are chosen by independent fair coin-flips; each of 15 players will get to see the colors of all the

[3]readers can find the problem in T. Ebert and H. Vollmer, "On the Autoreducibility of Random Sequences," http://www-info4.informatik.uni-wuerzburg.de/person/mitarbeiter/vollmer/publications.html.

other players' hats, and has the *option* of guessing the color of his own hat. There is to be no communication between the players; in particular, no player can tell what color a teammate has guessed or even whether he has guessed at all.

The players conspire beforehand and must come up with a strategy which maximizes the probability that *every* guess is correct, subject to the condition that it must guarantee that at least one player guesses.

As usual, there is an elegant solution and proof of optimality. Hint: the players can attain a 50% chance of all-correct by appointing one player to guess and the rest to pass. It's hard to believe, on first sight, that they can do any better; but in fact they can beat 90%!

Acknowledgment

I am obviously indebted to the (mostly) unknown inventors of the games described above, not to mention the many other wonderful mathematical brainteasers that have come my way over the years. If you are the originator, or know the originator, of any of them, I will be most grateful for a communication.

References

[1] T. Cover, Universal gambling schemes and the complexity measures of Kolmogorov and Chaitin, Statistics Department Technical Report #12, Stanford University, October 1974.

[2] T. Cover, Pick the largest number, in *Open Problems in Communication and Computation*, T. Cover and B. Gopinath, eds., Springer Verlag (1987), p. 152.

[3] T. Cover and J. Thomas, *Elements of Information Theory*, Wiley (1991), 132–133.

[4] K.S. Kaminsky, E.M. Luks and P.I. Nelson, Strategy, nontransitive dominance and the exponential distribution, *Austral. J. Statist.* **26** #2 (1984), 111–118.