Which is the Larger Number?

Sasha Gnedin*

November 29, 2012

Alice and Bob play a guessing game. Bob takes two identical cards and on each card writes a different number. The cards are shuffled and dealt face down on top of a table. Alice turns over a card and peeks at the number. She must then decide which of the two numbers is larger, winning a pound from Bob if her guess is right and losing a pound to Bob if it is wrong.

Guessing blindly, without any account of the revealed number, Alice will be right with probability one-half. Surprisingly, she has a decision strategy that performs better whichever the numbers. Here it is. Draw a threshold value t from the Gaussian distribution. If the revealed number exceeds t decide that the number is the larger of the two numbers on the cards, otherwise decide that the hidden number is the larger. For any two distinct numbers there is a nonzero probability p that the threshold falls between them, in which case the strategy returns the larger number as the only one exceeding the threshold. In the case the threshold falls below both numbers the strategy returns the revealed number, which by symmetry is equally likely to be the larger or the smaller. Likewise, in the case the threshold falls above both numbers the strategy returns the unrevealed number, which again is equally likely to be the larger or the smaller. Putting all three cases together, the total probability of correct guess is

$$p + \frac{1}{2}(1-p) = \frac{1}{2} + \frac{p}{2} > \frac{1}{2}.$$

It is known, and also detailed in this article, that despite Alice's apparent advantage the game is fair, in nice agreement with intuition. Still, the situation is paradoxical in two respects. Firstly, the blind guessing is a

*a.gnedin@qmul.ac.uk

minimax-optimal strategy, even though it is strictly dominated by another strategy (Gaussian). Secondly, Bob has no strategy to guarantee winning with probability exactly one-half, but playing smartly in an infinite series of games he can secure winning exactly half of the time in the long run. The bizarre features of the guessing game do not clash with the classic Minimax Theorem, which applies when both players have to choose from finitely many options.

We will not attempt explaining *why* the paradox occurs, rather present it as a fragment of a wider picture. To give the gist of it suppose Bob writes different numbers on 52 cards. The cards are shuffled and dealt face down into two piles of agreed sizes. Alice collects one of the piles and turns its cards over. She should decide which is the larger number – the largest of the revealed numbers or the largest of the remaining numbers hidden. If the piles are of equal size, a strategy based on comparing the maximum of the revealed numbers with the Gaussian threshold outperforms the blind guessing. If the piles are of distinct sizes, with, say, 25 and 27 cards, the situation is radically different. By blind guessing Alice can be right with probability 27/52. But now, in contrast to the symmetric case, Bob has a strategy, such that even Alice finds it out she cannot improve upon the blind guessing by benefiting from learning the numbers on the collected cards.

Two cards

If you offer the Gaussian strategy as a lunchtime brainteaser, ask your companions what would make a good strategy for Bob. You are likely to get 'Choosing the numbers close to one another, because the Gaussian p will be small'. No doubt that choosing big positive numbers, possibly with big difference, makes p small too. A missed point here is that a strategy of a player is rated by the outcome when the opponent finds the strategy out and reacts with a 'best response', which is a counter-strategy most unfavourable for the player. Within this basic paradigm of the zero-sum game theory, a *pure* strategy employing any fixed numbers is the worst Bob's behavior, because there exists a sure-fire guessing counter-strategy. Bob is better off using some randomization device to produce a probability distribution over the pairs of distinct numbers. Since the cards are shuffled anyway we may identify such a *mixed* strategy of Bob with a pair X, Y of random variables, that have a symmetry property called *exchangeability*, meaning that the joint probability distribution of X, Y is the same at that of Y, X. We shall agree that the revealed number is the value of X, for by exchangeability it does not matter how Alice picks a card to turn over.

A pure strategy of Alice can be regarded as a subset A of the real line: the strategy returns the revealed number as the larger if and only if X assumes a value in A. In particular, if A is the whole line Alice always decides on X, and if A is the empty set always on Y; these two are the *blind* strategies. A pure strategy is of the threshold type if A is a halfline $[t, \infty)$. The Gaussian strategy is a mixed strategy of the threshold type.

With the formal framework set out, it is useful to notice that strictly increasing functions $z \mapsto F(z)$ act on strategies of the players without altering outcome of the game, simply because $x > y$ is the same as $F(x) > F(y)$. In particular, any linear function $F(z) = \alpha z + \beta$ ($\alpha > 0$) maps a strategy in a strategy of the same quality.

The obvious randomization which comes to mind is drawing the numbers independently from some distribution, for instance with the aid of a random numbers generator. In the view of invariance under increasing F's, the specific shape of distribution does not matter, above that the distribution should be continuous, in order to have the probability of a tie $P(X = Y)$ equal to zero. Alice's best response is $A = [\mu, \infty)$, where the threshold is a *median* value for which $P(Y \leq \mu) = 1/2$. Arguing like for the Gaussian strategy, it is readily seen that the probability that X and Y fall on different sides from μ is $1/2$, hence Alice wins with probability $3/4$.

For exchangeable X, Y a best-response guessing strategy need not be of the threshold type. This is somewhat counter-intuitive, as it seems that the bigger the value of X the more likely it is the maximum of two numbers. A counter-example is a mixture of two pure strategies: one involving the numbers $1, 2$ and another the numbers $3, 4$. For instance, Bob can flip a coin to choose between the pairs. The best-response strategy of Alice is to decide that X is the larger number for $X = 2$, but not for $X = 3$.

Treating the question of the best-response guessing strategy in full generality, let X and Y be two arbitrary random variables (exchangeable or not), whose joint probability distribution is known and satisfies $P(X = Y) = 0$. With the value of X observed, a strategy maximizing the probability of guessing which of the numbers is the larger requires comparing the conditional probabilities. Given $X = x$ it is optimal to decide that x is the larger number if $P(Y < x|X = x) > 1/2$, and that x is the smaller number if

$P(Y < x|X = x) < 1/2$. If it occurs that

$$P(Y < x|X = x) = 1/2, \qquad (\spadesuit)$$

then it does not matter which decision to make.

For X, Y with given joint distribution, the probability of correct guess with the above strategy is strictly greater than one-half unless ($\spadesuit$) holds *identically* over the range of values of X. More precisely, the condition is that $P(Y < X|X) = 1/2$ should hold with probability one. Now, it turns that for exchangeable X, Y the identity is impossible, for otherwise the Gaussian strategy could not win with probability strictly greater than one-half!

Strategy S_n

If we ignore for a while the exchangeability, ($\spadesuit$) can be readily satisfied. Just take any integer random variable X, and define the conditional law of Y by $P(Y = x + 1|X = x) = P(Y = x - 1|X = x) = 1/2$ for every x. The unconditional distribution of Y is then computable as

$$P(Y = y) = \frac{1}{2}P(X = y - 1) + \frac{1}{2}P(X = y + 1).$$

If the exchangeability were to hold, we could replace Y by X, thus obtaining a linear recursion whose positive solutions are constants. The latter only makes sense as 'improper' uniform distribution on the infinite set of integers, with infinite total mass.

Let us approximate the 'improper' uniform by a proper distribution. To that end, it will be convenient to use another kind of two-step construction, which immediately implies exchangeability. First draw β from the uniform distribution over the set $\{1, \dots, n\}$, then choose between the pairs of values $\beta, \beta + 1$ and $\beta + 1, \beta$ for X, Y by tossing a fair coin. The resulting X has probability distribution with $n + 1$ positive weights

$$P(X = 1) = P(X = n + 1) = \frac{1}{2n}; \quad P(X = x) = \frac{1}{n} \qquad x \in \{2, \dots, n\},$$

which differ from the uniformity only on the edges. The edge values $X = 1$ and $X = n + 1$ are coupled with, respectively, $Y = 2$ and $Y = n$. But when $X = x$ for x any of $2, \dots, n$, the smaller number (which is β) could be $x - 1$ (then $Y = x - 1$) or x (then $Y = x + 1$), with each possibility being equally likely. Thus ($\spadesuit$) holds for all but the edge values of X.

We shall denote this Bob's strategy S_n (where 'S' stays for 'simplest'). Playing optimally against S_n Alice will be right with probability

$$P(X \in \{1, n+1\}) + \frac{1}{2} P(X \in \{2, \dots, n\}) = \frac{1}{n} + \frac{n-1}{2n} = \frac{1}{2} + \frac{1}{2n}.$$

Indeed, she wins if she reveals one of the edge numbers, and breaks even otherwise.

Fairness

Alice can secure winning with probability one-half. On the other hand, if Bob plays S_n he wins with probability as close as desired to one-half, provided n is large enough. Therefore the game is fair. With one-pound stakes the value of the game is zero.

Every threshold strategy of Alice, as well as both blind strategies are minimax. Bob has no minimax strategy, hence the game has no saddle point solution. The blind strategies are strictly dominated by any mixed threshold strategy, with threshold t sampled from a distribution with everywhere positive density.

Altering randomization from round to round, Bob can ensure that the limit proportion of the games won is exactly one-half. A composite strategy which neatly serves this purpose is playing S_{n^2} in round n: then in an infinite series of games Alice is kept breaking even every round, with just a few exceptions when X assumes an edge value 1 or n^2+1.

A model suitable to accommodate the nonstationary mixed strategies is a *repeated* game, defined as a sequence of identical stage games. Actually, admitting mixed strategies in a single game already presumes the infinite repetition, as it is inherent to the interpretation of probability in terms of relative frequencies. Assuming that players aim to maximize the limit proportion of the games won, the bonus of the extended framework is that the repeated game of guessing the larger number has a solution, and so the anomalies of a single game do not carry over.

Two piles

The game becomes more puzzling when the choice is to be made between two piles of cards. Suppose Bob takes k identical cards and on each writes a different number. The cards are turned face down, shuffled and dealt in two piles with ℓ and r cards ($k = \ell + r$), respectively. Alice gets the pile with ℓ cards and turns the cards over. She must then decide which of the piles has the card with the number largest of all k numbers. Guessing blindly on the pile with more cards Alice is right with probability $\rho = \max(\ell/k,\ r/k)$.

With the account of shuffling, a mixed strategy of Bob can be regarded as a sequence $Z_1, \dots, Z_k$ of exchangeable random variables with pairwise distinct values. Exchangeability means that the joint distribution of the sequence is not changed by $k!$ permutations of the variables. Let $X = \max(Z_1, \dots, Z_\ell)$ be the largest number in the collected pile, and let $Y = \max(Z_{\ell+1}, \dots, Z_k)$ be the largest in the remaining pile. Having learned the values of $Z_1, \dots, Z_\ell$, Alice is willing to guess which of X and Y is the larger number.

If $\ell = r$ the piles maxima X and Y are exchangeable, hence the blind strategy is dominated, like in the game with two cards. Nonetheless, finding a good strategy for Bob and showing fairness is by far not straightforward, unless $\ell = r = 1$. The asymmetric case $\ell \neq r$ brings new challenges. The intuition suggests that it is possible to benefit from learning the numbers, especially when it is known which random rule has produced them. This fails, however, as we about to show.

The largest in the pile

Alice's pure strategy is a set A of ℓ-tuples of distinct numbers, such that A is preserved by permutations of the coordinates. In particular, for given joint distribution of $Z_1, \dots, Z_k$, Alice's best-response strategy is to decide that $X = x$ is the overall largest number if and only if

$$P(Y < x | Z_1 = z_1, \dots, Z_\ell = z_\ell) \geq 1/2\,,$$

where $x = \max(z_1, \dots, z_\ell)$.

The inequality is easy to deal with if $Z_1, \dots, Z_k$ are independent and identically distributed, because then X and Y independent. Assuming Z_j's uniformly distributed on $[0, 1]$, the criterion for guessing on X becomes $x \geq 2^{-1/r}$, where $2^{-1/r}$ is the median value of Y. With some effort Alice's winning probability is computed as $(\ell + r2^{-\ell/r})/(\ell + r)$. The probability is the same for any continuous distribution. The formula gives $3/4$ for $\ell = r = 1$, in agreement with our previous findings.

If X and Y are not independent the required conditional probability depends $z_1, \dots, z_\ell$ in a more complex way. Fortunately, for our purposes it is possible to focus on a special class of exchangeable distributions, for which only $x = \max(z_1, \dots, z_\ell)$ matters, like in the independent case, or like in the game with two cards.

Consider positive random variables $Z_1, \dots, Z_k$ with the joint density of the form

$$f_k(z_1, \dots, z_k) = g_k(\max(z_1, \dots, z_k)). \qquad (\diamond)$$

The exchangeability holds, because max is a symmetric function. To define a density, the function g_k must be nonnegative and satisfy the integrability condition

$$\begin{aligned} \int_0^\infty \dots \int_0^\infty f_k(z_1, \dots, z_k) dz_1 \dots dz_k &= \\ k \int_0^\infty g_k(z_1) dz_1 \int_0^{z_1} \dots \int_0^{z_1} dz_2 \dots dz_k &= \\ k \int_0^\infty g_k(t) t^{k-1} dt &= 1. \end{aligned}$$

The joint density of revealed numbers $Z_1, \dots, Z_\ell$ is obtained by integrating out r variables

$$f_\ell(z_1, \dots, z_\ell) = \int_0^\infty \dots \int_0^\infty g_k(\max(x, z_{\ell+1}, \dots, z_k)) dz_{\ell+1} \dots dz_k = \\ x^r g_k(x) + r \int_x^\infty y^{r-1} g_k(y) dy =: g_\ell(x),$$

which is again a density depending only on the maximum of the arguments. The conditional probability we are interested in depends only on x,

$$P(Y < x | Z_1 = z_1, \dots, Z_\ell = z_\ell) = P(Y < x | X = x) = \frac{x^r g_k(x)}{g_\ell(x)}, \qquad (\heartsuit)$$

as promised.

Bob's strategy

The ideal choice of g_k is the function $g_k(x) = x^{-k}$, which turns the right-hand side of ($\heartsuit$) in constant ℓ/k. This is too good to be true, and, indeed, this function does not meet the integrability condition, so does not define a density f_k. As a useable approximation to the ideal, we take a piece-wise smooth function

$$g_k(x) = \begin{cases} cx^{-k+\delta}, & 0 < x < 1, \\ cx^{-k-\delta}, & x \geq 1, \end{cases} \tag{$\clubsuit$}$$

where $c = \delta/(2k)$ is the normalization constant, and $\delta > 0$ is yet to be determined.

For $x \geq 1$ we readily compute

$$g_\ell(x) = cx^{-\ell-\delta}\left(1 + \frac{r}{\ell+\delta}\right),$$

which substituted in the right-hand side of ($\heartsuit$) yields a constant

$$P(Y < x|X = x) = \frac{\ell+\delta}{k+\delta}.$$

For $0 < x < 1$ a more involved formula emerges

$$P(Y < x|X = x) = \frac{\ell-\delta}{k-\delta-2\delta r x^{\ell-\delta}/(\ell+\delta)}.$$

Putting the two cases together, we see that for every $\epsilon > 0$, choosing δ small enough we can achieve that the inequality

$$|P(Y < x|X = x) - \ell/k| < \epsilon$$

holds for all $x > 0$ and $\ell = 1, \dots, k$.

Manipulating ϵ, in the case $\ell < r$ we achieve that $P(Y < x|X = x) < 1/2$ holds for all $x > 0$, implying that the best-response strategy of Alice is to always decide on the larger pile with r cards, no matter what the revealed numbers. For $\ell > r$ the inequality is reversed, hence it is again optimal to always decide on the larger pile, which in this case is the collected pile with ℓ cards. We see that for $\ell \neq r$ Bob has means to force Alice guessing blindly on the pile with more cards, thereby keeping her winning probability by at most ρ.

It follows that both players have minimax strategies, and the game has a saddle-point solution. The blind strategy of Alice is the unique minimax strategy. Bob has many minimax strategies: there is a freedom in selecting δ, and more strategies can be obtained, for instance, by applying increasing functions. For $\ell = r$ the constructed distributions are still good, and Bob can use one with small enough δ to win with probability as close to $1/2$ as desired.

Bob need not even care about the size of the revealed pile, playing the same strategy for every $\ell = 1, \dots, k-1$.

The Bayesian viewpoint

Bob's strategy has a useful representation, which enables a simple implementation by means of a standard uniform random numbers generator. First, obtain $u_0, u_1, \dots, u_k$ using the generator. Then find α by solving the quantile equation

$$\int_0^\alpha \varphi(t)dt = u_0,$$

where

$$\varphi(\alpha) = \begin{cases} \frac{\delta}{2}(1 - \frac{\delta}{k})\alpha^{\delta-1}, & 0 < \alpha < 1, \\ \frac{\delta}{2}(1 + \frac{\delta}{k})\alpha^{-\delta-1}, & \alpha \geq 1 \end{cases}$$

is a density of the positive half-line. It is not hard to find α explicitly. Finally, take $z_1 = \alpha u_1, \dots, z_k = \alpha u_k$ as realization of $Z_1, \dots, Z_k$.

To see why this works, note that for $U_1, \dots, U_k$ independent, uniformly distributed on $[0, 1]$, the random variables $\alpha U_1, \dots, \alpha U_k$ are independent uniform on $[0, \alpha]$, thus with constant joint density α^{-k} in the k-dimensional cube of side α. For φ a density on the positive half-line, sampling the factor α from φ (independently from U_j's) yields a mixed k-dimensional density

$$f_k(z_1, \dots, z_k) = \int_{\max(z_1,\dots,z_k)}^\infty \alpha^{-k}\varphi(\alpha)d\alpha,$$

which is of the form ($\diamond$). Sampling α can be conducted by solving the quantile equation with yet another uniform U_0 – this is a well-known trick to generate a nonuniform random number. Specializing φ as above results in ($\diamond$) with g_k given by ($\clubsuit$).

If Alice finds out Bob's strategy, she faces a Bayesian decision problem. Then she knows in advance that the numbers will be drawn independently from the uniform distribution on $[0, \alpha]$, where the parameter α has prior density φ. (If she knew α exactly, she could improve upon the blind strategy by guessing like in the independent case.) After the first ℓ numbers have been observed, her knowledge of the parameter is summarized as a posterior density φ_ℓ, which relates to g_ℓ like φ relates to g_k. When ℓ is large, the update seem to give a good idea of what the value of α actually is, but this information turns insufficient to help guessing.

A general principle of statistical decision theory suggests that if elements of a problem are invariant under the action of a group of transformations, then the procedures adopted are also invariant. Uniform distribution on $[0, \alpha]$ ($\alpha > 0$) is a scale family, derived from the standard uniform distribution by monotone transformations $x \mapsto \alpha x$. Composition of the transformations corresponds to the multiplication on the parameter space $(0, \infty)$, and the measure with density k/α is an invariant measure on the group. Using the measure as 'improper prior' leads to a 'solution' that corresponds to the ideal function $g_k(x) = x^{-k}$.

In a similar line, continuous analogues of S_n can be associated with location families of distributions. For instance, a good strategy for Bob in the game with two cards is to draw β from the uniform distribution on a large interval, then draw two numbers independently from the uniform distribution on $[\beta - 1/2, \beta + 1/2]$.

Comments

A randomized threshold strategy beating the blind guessing originated in a half-page abstract of the information theorist Tom Cover [3]. Precursors are found in the work on theoretical statistics [1, 5]. The idea has been popularized since then in [2, 6, 14]. Threshold strategies were exploited in [5, 10] to show that $P(Z_k = \max(Z_1, \ldots, Z_k) | Z_1, \ldots, Z_{k-1}) = 1/k$ and some other analogues of (♠) cannot hold for exchangeable Z_j's without ties. Strategy S_n for the number writer was outlined by Ted Hill [6] and Peter Winkler [14]. In [12] the game with two cards is connected to the famous two-envelopes paradox.

Sequential versions of the game with k cards, related to the famous *secretary problem* and Martin Gardner's *game of googol*, were intensively discussed in [4, 7, 9, 13]. The use of scale mixtures of uniform distributions in this context has some history, for which we refer to [4, 8]. It is known [11] that the location mixtures of uniform distributions for large k are far from optimality for the number writer. For $k > 2$ no good strategies with integer numbers are known explicitly.

References

[1] D. Blackwell (1951), On the translation parameter problem for discrete variable, *Annals of Mathematical Statistics* 22: 391–399.

[2] F.T. Bruss (2003), Playing a trick on uncertainty, *Newsletter of the European Mathematical Society*, Issue 50, pp. 5 and 7-8 (online).

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

[4] T. Ferguson (1989), Who Solved the Secretary Problem? *Statistical Science* 4: 282–289.

[5] B.M. Hill (1968), Posterior distribution of percentiles: Bayes' theorem for sampling from a population, *Journal of the American Statistical Association* 63: 677–691.

[6] T.P. Hill (2009) Knowing when to stop, *American Scientist* 97 No.2 (online).

[7] T.P. Hill and U. Krengel (1992), On the game of googol, *International Journal of Game Theory* 21: 151–160.

[8] A. Gnedin (1995), A solution to the game of googol, *Annals of Probability* 22: 1588–1595.

[9] A. Gnedin and U. Krengel (1995), A stochastic game of optimal stopping and order selection, *Annals of Applied Probability* 5: 310–321.

[10] A. Gnedin and U. Krengel (1996), Optimal selection problems based on exchangeable trials, *Annals of Applied Probability* 6: 862–882.

[11] J. Petruccelli (1980), On a best choice problem with partial information, *Annals of Statistics* 8: 1171–1174.

[12] D. Samet, I. Samet and D. Schmeidler (2004), One observation behind two-envelope problems, *American Mathematical Monthly* 111:347–351.

[13] S.M. Samuels (1989), Who will solve the secretary problem? [Comment on Ferguson] *Statistical Science* 4: 289–291.

[14] P. Winkler (2001), Games people don't play, in *Puzzlers Tribute: A Feast for the Mind*, D. Wolfe and T. Rodgers eds., AK Peters.