

Knowing When to Stop

How to gamble if you must—the mathematics of optimal stopping

Theodore Hill

Every decision is risky business. Selecting the best time to stop and act is crucial. When Microsoft prepares to introduce Word 2020, it must decide when to quit debugging and launch the product. When a hurricane veers toward Florida, the governor must call when it's time to stop watching and start evacuating. Bad timing can be ruinous. Napoleon learned that the hard way after invading Russia. We face smaller-consequence stopping decisions all the time, when hunting for a better parking space, responding to a job offer or scheduling retirement.

The basic framework of all these problems is the same: A decision maker observes a process evolving in time that involves some randomness. Based only on what is known, he or she must make a decision on how to maximize reward or minimize cost. In some cases, little is known about what's coming. In other cases, information is abundant. In either scenario, no one predicts the future with full certainty. Fortunately, the powers of probability sometimes improve the odds of making a good choice.

While much of mathematics has roots that reach back millennia to Euclid and even earlier thinkers, the history of probability is far shorter. And its lineage is, well, a lot less refined. Girolamo Cardano's famed 1564 manuscript *De Ludo Aleae*, one of the earliest writings on probability and not published until a century after he wrote it, primarily analyzed dice games. Although Galileo and other 17th-century scientists contributed to this enterprise, many credit the mathematical foundations of probability to an exchange of letters in 1654 between two famous French mathematicians, Blaise Pascal and Pierre de Fermat. They too were concerned with odds and dice throws—for example, whether it is wise to bet even money that a pair of sixes will occur in 24 rolls of two fair dice. Some insisted it was, but the true probability of a double six in 24 rolls is about 49.1 percent.

That correspondence inspired significant advances by Abraham de Moivre, Christiaan Huygens, Siméon Poisson, Jacob Bernoulli, Pierre-Simon Laplace and Karl Friedrich Gauss into the 19th century. Still, for a long time there was no formal definition of probability precise enough for use in mathematics or robust enough to handle increasing evidence of random phenomena in science. Not until 1933, nearly four centuries after Cardano, did the Russian mathematician Andrey Kolmogorov put probability theory on a formal axiomatic basis, using the emerging field we now call measure theory.

The history of optimal-stopping problems, a subfield of probability theory, also begins with gambling. One of the earliest discoveries is credited to the eminent English mathematician Arthur Cayley of the University of Cambridge. In 1875, he found an optimal stopping strategy for purchasing lottery tickets. The wider practical applications became apparent gradually. During World War II, Abraham Wald and other mathematicians developed the field of statistical sequential analysis to aid military and industrial decision makers faced with strategic gambles involving massive amounts of men and material. Shortly after the war, Richard Bellman, an applied mathematician, invented dynamic programming to obtain optimal strategies for many other stopping problems. In the 1970s, the theory of optimal stopping emerged as a major tool in finance when Fischer Black and Myron Scholes discovered a pioneering formula for valuing stock options. That transformed the world's financial markets and won Scholes and colleague Robert Merton the 1997 Nobel Prize in Economics. (Black had died by then.)

The Black-Scholes formula is still the key to modern option pricing, and the optimal-stopping tools underlying it remain a vigorous area of research in academia and industry. But even elementary tools in the theory of optimal stopping offer powerful, practical and sometimes surprising solutions.

The Marriage Problem

Suppose you decide to marry, and to select your life partner you will interview at most 100 candidate spouses. The interviews are arranged in random order, and you have no information about candidates you haven't yet spoken to. After each interview you must either marry that person or forever lose the chance to do so. If you have not married after interviewing candidate 99, you must marry candidate 100. Your objective, of course, is to marry the absolute best candidate of the lot. But how?

This problem has a long and rich history in the mathematics literature, where it is known variously as the marriage, secretary, dowry or best-choice problem. Certainly you can select the very best spouse with probability at least 1/100, simply by marrying the first person. But can you do better? In fact, there is a simple rule that will guarantee you will marry the absolute best more than one-third of the time. And the rule can be transferred to other scenarios.



[+ enlarge image](#)



[+ enlarge image](#)



enlarge image

As an enlisted man in the U.S. Air Force during the Vietnam era, John Elton, now a Georgia Institute of Technology mathematician, transformed the marriage problem into a barracks moneymaking scheme. Elton asked his fellow airmen to write down 100 different numbers, positive or negative, as large or small as they wished, on 100 slips of paper, turn them face down on a table and mix them up. He would bet them he could turn the slips of paper over one at a time and stop with the highest number. He convinced them it was "obvious" that the chance of him winning was very small, so he asked ten dollars if he won and paid them one dollar if he lost. There was no shortage of bettors. Even though my friend lost nearly two-thirds of the games, he won more than one-third of them. And with the 10-1 odds, he raked in a bundle. How?

First, note that there is a very simple strategy for winning more than one-fourth of the time, which would already put him ahead. Call an observed number a "record" if it is the highest number seen so far. Suppose you turn over half the numbers—or interview the first 50 marriage candidates—and never stop, no matter how high the number. After that you stop with the first record you see. If the second-highest number in the 100 cards happens to be in the first 50 you look at, and the highest in the second half—which happens 1 in 4 times—then you win. That strategy is good, but there is an even better one. Observe only 37 cards (or potential partners) without stopping and then stop with the next record. John Gilbert and Frederick Mosteller of Harvard University proved that this strategy is best and guarantees stopping with the best number about 37 percent of the time. In fact, observing $N/e = 0.37$ of the candidates, where N is the total number of candidates and e is the base of the natural logarithms, $e = 2.71828\dots$, guarantees winning with probability more than $1/e > 0.36$, no matter how many cards or candidates there are. (Note that the "observe half the numbers" strategy clearly wins with probability at least $1/4$, also independent of the number of cards.)

Sometimes the goal is to stop with one of the best k of N candidates. That is, you win if you stop with any one of the highest k numbers. In the Olympics or in horse racing, for example, the objective often is the $k = 3$ case—to win a medal or to show—rather than the all-or-nothing $k = 1$ goal of a gold medal or a win, which is much riskier. The optimal strategy for stopping with one of the best k is similar to stopping with the best. First, you should observe a fixed number of candidates without ever stopping, thereby obtaining a baseline to work with. Then for another certain fixed length of time, stop if you see a record. Since it is the best seen so far, it is somewhat likely to be one of the best k . If no record appears during that stretch, then continue to the next stage where you stop with one of the highest two numbers for a fixed period of time, and so on. For $k = 2$, this method guarantees a better than 57 percent chance of stopping with one of the two best even if there are a million cards. For small N , the probability is quite high. Figure 3 illustrates the optimal strategy for $N = 7$ and $k = 2$, which guarantees a win two-thirds of the time.

$N = 2$ Surprise

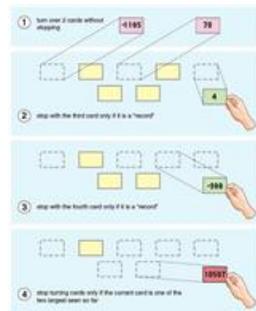
Now, suppose you must decide when to stop and choose between only two slips of paper or two cards. You turn one over, observe a number there and then must judge whether it is larger than the hidden number on the second. The surprising claim, originating with David Blackwell of the University of California, Berkeley, is that you can win at this game more than half the time. Obviously you can win exactly half the time by always stopping with the first number, or always stopping with the second, without even peeking. But to win more than half the time, you must find a way to use information from the first number to decide whether or not to stop. (Readers take comfort: When mathematicians first heard this claim, many of us found it implausible.)

Here is one stopping rule that guarantees winning more than half the time. First, generate a random number R according to a standard Gaussian (bell-shaped) curve by using a computer or other device. Then turn over one of the slips of paper and observe its number. If R is larger than the observed number, continue and turn over the second card. If R is smaller, quit with the number observed on the first card. How can such a simple-minded strategy guarantee a win more than half the time?

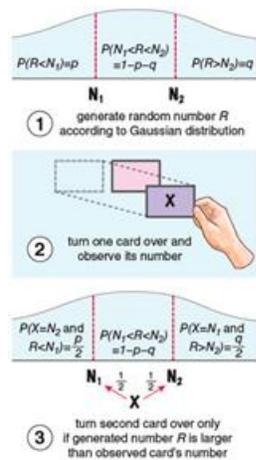
If R is smaller than each of the two written numbers, then you win exactly half the time ($p / 2$ of the unknown probability p in Figure 4); if it is larger than both, you again win half that time ($q / 2$ of q , also in Figure 4). But if R falls between the two written numbers, which it must do with strictly positive probability (since the two numbers are different and the Gaussian distribution assigns positive probability to every interval) then you win *all* the time. This gives you the edge you need, since $p / 2 + q / 2 + 1 - p - q$ is greater than $1/2$, because $1 - p - q$ is greater than zero. For example, if the two hidden numbers are 1 and n , this Gaussian method yields a value for p about .8413 and q about .0008, so the probability that it will select the larger number is more than 57 percent.

Of course if the number writer knows this Gaussian strategy, he can make your winnings as close to $1/2$ as he wants by writing numbers that are very close.

If the number writer is not completely free to pick any number, but instead is required to choose an integer in the range $\{1, 2, \dots, 100\}$, say, then he cannot make your probability of winning arbitrarily close to $1/2$. In this case it also seems obvious that the number-writer would never write a 1, since if you turn over a 1, you will always win by not stopping. But if he never writes a 1, he then would never write a 2 either since he never wrote a 1, and so on *ad absurdum*. Interested readers are invited to discover for themselves the optimal strategy in this case, and the amount more than $1/2$ one can guarantee to win on the average.



enlarge image



enlarge image

Backward Induction

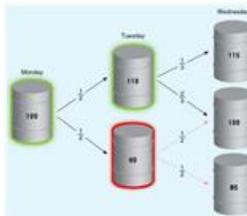
At the opposite end from having no information about future values is having full information—that is, complete information about the probabilities and the exact values of all potential future observations. In the spirit of Cardano, Fermat and Pascal's discoveries about probability with dice centuries ago, let's consider a game of rolling a standard, fair, six-sided die at most five times. You may stop whenever you want and receive as a reward the number of Krugerrands corresponding to the number of dots shown on the die at the time you stop. (At the time of this writing, one Krugerrand is worth \$853.87.) Unlike the no-information marriage problems, here everything is known. The values at each roll will be 1, 2, 3, 4, 5, or 6, and the probability of each number on each roll is one-sixth. The objective is to find the stopping rule that will maximize the number of Krugerrands you can expect to win on average.

If you always stop with the first roll, for example, the winnable amount is simply the expected value of a random variable that takes the values 1, 2, 3, 4, 5, and 6 with probability $1/6$ each. That is, one-sixth of the time you will win 1, one-sixth of the time you will win 2, and so on, which yields the expected value $1(1/6) + 2(1/6) + 3(1/6) + 4(1/6) + 5(1/6) + 6(1/6) = 7/2$. Thus if you always quit on the first roll, you expect to win 3.5 Krugerrands on the average. But clearly it is not optimal to stop on the first roll if it is a 1, and it is always optimal to stop with a

total number of rolls	stop if initial roll is:	average optimal expected reward
1	{1, 2, 3, 4, 5, 6}	3.5
2	{4, 5, 6}	4.25
3	{5, 6}	4.67
4	{5, 6}	4.94
5	{5, 6}	5.13

enlarge image

6, so already you know part of the optimal stopping rule. Should you stop with a 5 on the first roll? One powerful general technique for solving this type of problem is the method of *backward induction*.



+ enlarge image

Clearly it is optimal to stop on the first roll if the value seen on the first roll is greater than the amount expected if you do not stop—that is, if you continue to roll after rejecting the first roll. That would put you in a new game where you are only allowed four rolls, the expected value of which is also unknown at the outset. The optimal strategy in a four-roll problem, in turn, is to stop at the first roll if that value is greater than the amount you expect to win if you continue in a three-roll problem, and so on. Working down, you arrive at one strategy that you do know. In a one-roll problem there is only one strategy, namely to stop, and the expected reward is the expected value of one roll of a fair die, which we saw is 3.5. That information now yields the optimal strategy in a two-roll problem—stop on the first roll if the value is more than you expect to win if you continue, that is, more than 3.5. So now we know the optimal strategy for a two-roll problem—stop at the first roll if it is a 4, 5, or 6, and otherwise continue—and that allows us to calculate the expected reward of the strategy.

In a two-roll problem, you win 4, 5, or 6 on the very first roll, with probability $1/6$ each, and stop. Otherwise (the half the time that the first roll was a 1, 2 or 3) you continue, in which case you expect to win 3.5 on the average. Thus the expected reward for the two-roll problem is $4(1/6) + 5(1/6) + 6(1/6) + (1/2)(3.5) = 4.25$. This now gives you the optimal strategy for a three-roll problem—namely, stop if the first roll is a 5 or 6 (that is, more than 4.25), otherwise continue and stop only if the second roll is a 4, 5, or 6, and otherwise proceed with the final third roll. Knowing this expected reward for three rolls in turn yields the optimal strategy for a four-roll problem, and so forth. Working backwards, this yields the optimal strategy in the original five-roll problem: Stop on the first roll only if it is a 5 or 6, stop on the second roll if it is a 5 or 6, on the third roll if it is a 5 or 6, the fourth roll if it is a 4, 5 or 6, and otherwise continue to the last roll. This strategy guarantees that you will win about 5.12 Krugerrands on average, and no other strategy is better. (So, in a six-roll game you should stop with the initial roll only if it is a 6.)

The method of backward induction is very versatile, and works equally well if the process values are not independent (as they were assumed to be in repeated rolls of the die), or if the objective is to minimize some expected value such as cost. Suppose that a company must purchase its weekly energy supplies on the previous Monday, Tuesday or Wednesday and the likelihood of future prices can be estimated based on past statistics. For example, if on Monday the decision maker has the opportunity to purchase energy for the following week at a cost of 100, she may know from past experience that there is a 50-50 chance that Tuesday's price will be 110, and otherwise it will be 90. Furthermore she knows that if it is 110 on Tuesday, then it will be 115 on Wednesday with probability $1/3$, and otherwise will be 100; and that if it is 90 on Tuesday, it is equally likely to be 100 or 85 on Wednesday. Using backward induction, the optimal rule on Tuesday is seen to be not to buy if the price is 110, since 110 is larger than the expected price $(1/3)(115) + (2/3)(100) = 105$ if she delays buying until Wednesday. Similarly, if the price Tuesday is 90, it is optimal to buy. Working backwards to Monday, since 100 is larger than the expected cost if she continues—namely, $(1/2)(105) + (1/2)(90) = 97.5$ —it is optimal not to buy on Monday. Putting these together yields her optimal stopping strategy.

In the full-information case, with the objective of stopping with one of the largest k values, the best possible probabilities of winning were unknown for general finite sequences of independent random variables. Using both backward and forward induction, and a class of distributions called "Bernoulli pyramids," where each new variable is either the best or the worst seen thus far (for example, the first random variable is either +1 or -1 with certain probabilities, the second variable is +2 or -2, and so forth), Douglas Kennedy of the University of Cambridge and I discovered those optimal probabilities. We proved that for every finite sequence of independent random variables, there is always a stop rule that stops with the highest value with probability at least $1/e$, and a stop rule that stops with one of the two highest values with probability at least

$$e^{-\sqrt{2}}(1+\sqrt{2}) \cong 0.59,$$

+ enlarge image

which is the best you can do. (The probabilities for stopping with one of the k values highest have similar formulas).

Computers were not useful for solving that problem. In fact, all of the problems described in this article were solved using traditional mathematicians' tools—working example after example with paper and pencil; settling the case for two, three and then four unknowns; looking for patterns; waiting for the necessary *Aha!* insights; and then searching for formal proofs for each step. Computers are very helpful for after-the-fact applications of many results, such as backward induction. But in theoretical probability, computers often do not significantly aid the discovery process. To better understand this, reconsider the two-card problem described above. How would one program a computer to imagine such a strategy existed, let alone to search the universe of ideas to find it?

Partial Information

The case of partial information is the most difficult. Usually one does not know how many applicants there will be for a job, nor the exact probabilities of future stock values. In these situations, one method of solution is to use tools from the theory of zero-sum, two-person games, where the stopping problem can be thought of as a decision maker playing against an opponent (often called Nature or God) who may assign the values and probabilities in any way she chooses.

Ulrich Krengel of the University of Göttingen and I used this technique to discover the optimal strategy in the so-called marriage problem where only a bound on the number of applicants is known. As a concrete example, consider the problem where the objective is to select the highest number from a hat containing at least one, and at most five, numbered cards (if you do not stop and there are no cards left, you lose.) We proved that the optimal strategy in this case is to stop with the first card with probability $26/75$ (which you may do using a random number generator or other method). If you do not stop with the first card, then you should continue to the second card, if there is one. If the second card's number is higher than the first card's number, stop with probability $26/49$. Otherwise, continue, stopping with the first record thereafter (or when you run out of cards or are forced to choose the number on the fifth card). This guarantees a probability of $26/75$ of stopping with the highest number, no matter how many cards Nature deposited in the hat.

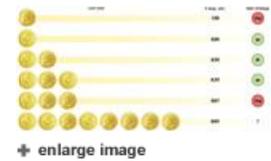
There is no better strategy. We found the exact formulas and strategies for all possible bounds on the maximum number of cards and the winning probabilities are surprisingly high. For example, even if you only know that there are somewhere between 1 and 100 cards in the hat, it is still possible to win about 20 percent of the time. Exactly the same method can be employed to obtain optimal stopping rules in many real-life problems such as a situation where an employer wants to hire the very best salesperson available, and knows the maximum number of candidates available for the position, but does not know how many have already accepted other jobs.

In another type of stopping problem involving partial information, the observer knows the length of the sequence exactly (say, for example, the number of cards), but has only partial information about the random values on the cards. Instead of having no information at all, or knowing all the possible values and probabilities, he might know only the average value and standard deviation of each variable. In the case where variables are independent, Frans Boshuizen of the Free University of Amsterdam and I were able to use game theory and mass-shifting "balayage" (sweeping probability mass away from its center in both directions) arguments to determine the optimal stop rules, but those techniques fail for most other partial-information stopping problems.

Unsolved Problems

Although many stopping problems have been solved, there are still tantalizingly simple unsolved problems, even ones

involving full information. My favorite is this: Toss a fair coin repeatedly and stop whenever you want, receiving as a reward the average number of heads accrued at the time you stop. If your first toss is a head, and you stop, your reward is 1 Krugerrand. Since you can never have more than 100 percent heads, it is clearly optimal to stop in that case. If the first toss is a tail, on the other hand, it is clearly best not to stop, since your reward would be zero. Suppose the first toss is a tail and the second a head. You may quit then and receive half a Krugerrand, or continue to toss. A moment's reflection shows that it is never optimal to stop with half a Krugerrand or less, since the law of large numbers says that the average number of heads will converge over time to 50 percent, and in doing so will oscillate above and below 50 percent repeatedly. Stopping with 50 percent or less is simply not greedy enough.



With a little more difficulty it can be shown that stopping with the third toss if you saw tail-head-head is optimal, and that stopping the very first time you observe more heads than tails is optimal for a while. But stopping the first time you have one more head than tails is not optimal forever. After a certain critical time you should only stop when you have two more heads than tails, and after a second critical time, stop only when you are three heads ahead, and so forth. The proof of this fact, which relies on the law of the iterated logarithm, is not easy, and the complete list of critical times is still not known. Backward induction will not work for this problem since there is no *a priori* end to the sequence and, hence, no future time to calculate backwards from. Even though Wolfgang Stadje of the University of Osnabrück has advanced this problem very recently, and despite the gains of a century of development of mathematical probability, the exact optimal rule for all sequences of heads and tails is unknown.

Still, the general field of optimal stopping, especially with its applications to financial markets, continues to develop at a rapid pace. In fact, some experts feel that the pace has been too quick and that computer models of option and derivative pricing (basic optimal-stopping problems) are the roots of the current economic crisis. But it is not the theory that is at fault. Like Steven Shreve of Carnegie Mellon University, I blame the decision makers' blind trust in computer model predictions. In fact, new ideas and discoveries in optimal stopping, including better estimates of the risk that mathematical models are wrong, are exactly what we need—not only for guidance on when to stop the bailouts, for example, but also for help with multiple other monumental problems, including when to stop using fossil fuels or stockpiling nuclear weapons.

Bibliography

- Boshuizen, F., and T. Hill. 1992. Moment-based minimax stopping functions for sequences of random variables. *Stochastic Processes and Applications* 43:303-316.
- Bruss, F. T. 2000. Sum the odds to one and stop. *Annals of Probability* 28:1384-1391.
- Chow, Y., H. Robbins and D. Siegmund. 1991. *Great Expectations: The Theory of Optimal Stopping*. Mineola, N.Y.: Dover Publications.
- Cover, T. 1987. Pick the largest number, in *Open Problems in Communication and Computation*, eds. Cover, T., and B. Gopinath. New York: Springer-Verlag, p. 152.
- Dubins, L., and L. Savage. 1976. *Inequalities for Stochastic Processes: How to Gamble If You Must*. New York: Dover Publications.
- Ferguson, T. 1989. Who solved the secretary problem? *Statistical Science* 4:282-296.
- Freeman, P. R. 1983. The secretary problem and its extensions: A review. *International Statistical Review* 51:189-208.
- Hill, T., and U. Krengel. 1992. Minimax-optimal stop rules and distributions in marriage problems. *Annals of Probability* 19:342-353.
- Hill, T., and D. Kennedy. 1992. Sharp inequalities for optimal stopping with rewards based on ranks. *Annals of Applied Probability* 2:503-517.
- Jacka, S., et al. 2007. Optimal stopping with applications. *Stochastics* 79: 1-4.
- Peskir, G., and Albert Shiryaev. 2006. *Optimal Stopping and Free Boundary Problems*. Boston: Birkhäuser.
- Samet, D., I. Samet and D. Schmeidler. 2004. One observation behind two-envelope problems. *American Mathematical Monthly* 111:347-351.
- Shreve, S. 2008. Don't blame the quants. *Forbes.com*. http://www.forbes.com/2008/10/07/securities-quants-models-oped-cx_ss_1008shreve.html
- Stadje, W. 2008. The maximum average gain in a sequence of Bernoulli trials. *American Mathematical Monthly* 115:902-910.

You can find this online at <http://www.americanscientist.org/issues/num2/2009/2/knowning-when-to-stop/1>
 © Sigma Xi, The Scientific Research Society