# The Maths of Game Theory

Professor Sarah Hart

22nd November 2022

## Introduction

When we buy, sell, bargain, barter, bid at auctions, and compete for resources, we want to be sure that we are using the best strategies. Mathematics, particularly the field known as game theory, can help us understand precisely these kinds of situations. That's why in 1994, the Nobel Prize for Economics was won by a mathematician – John Nash. Using games like the Prisoner's dilemma, this lecture explains the work of game theorists such as Nash, David Blackwell and John von Neumann. Along the way we'll find out how much to charge for a cup of coffee, the best way to fight a duel, and when is the right time to buy a house.

We'll begin with just a couple of examples of situations that can arise in real life around buying, selling and other economic activity. In "price war", the question is, if I'm running a company, should I lower my prices to increase market share – the lower price means less profit per item, but if I'm selling many more items I'm going to increase overall profit. So what should I do? In "best coffee in town", café owners at a tourist hotspot can charge very high prices to tourists who are never going to return anyway, and can't try all the cafés to find out where the cheapest and best coffee is, but if they charge too much, the locals will stay away. So what's the best price to set? Finally, in an auction, especially a sealed bid or silent auction, what price should you bid to avoid being disappointed? In all these cases, the decisions you make are affected by the decisions of others, just like in a game of chess. Analysing situations like these mathematically by thinking of them as games, with players and strategies, is what we call game theory. I'm going to show you how some simple mathematical ideas can help us understand "games" like these, which allows us both to find good strategies, but also to design the rules of the game to get a better outcome.

## Price Wars and the Prisoner's Dilemma

I'm going to begin with "price war". This is actually a version of a very famous problem known as the Prisoner's Dilemma. Lots of people have heard of this but I didn't feel I could give a talk involving game theory without mentioning it, so do bear with me if you are familiar with it already. In the Prisoner's Dilemma, two people commit a crime, say a robbery, and are arrested. But without more evidence the police can only hope for a conviction on the lesser charge of possessing stolen goods. If you implicate your accomplice, and your accomplice stays silent, then you go free and they get the full term, say 5 years. If you both talk, then you both get 3 years, as you are both guilty but you cooperated with the police. We can represent these outcomes in a grid.

| | Prisoner B talks | Prisoner B stays silent |
|---|---|---|
| Prisoner A talks | A gets 3 years<br>B gets 3 years | A gets 0 years<br>B gets 5 years |
| Prisoner A stays silent | A gets 5 years<br>B gets 0 years | A gets 1 year<br>B gets 1 year |

If you are Prisoner A, what should you do, if acting purely from self-interest? You don't know what B will do. If B talks, then looking at that column, it's clear the least worst outcome for you is to talk (3 years rather than 5). If B stays silent, then again the best choice is to talk, because then you go free rather than serving 1 year. So whatever B does, you should talk. The paradox is that B faces the same decision, and will reach the same conclusion. This means that both prisoners would talk, and so both serve 3 years, rather than the logically best outcome (in terms of total jail time served) of both staying silent. It's an instance where if the "players" co-operate, they can achieve a better outcome. Of course, in a situation like this there may be other factors at play. Implicating your accomplice is frowned upon in some circles.

Exactly the same dynamic is at play on our "price war" situation. Suppose Company A and Company B both make widgets. They cost £5 to make, and both companies currently sell them for £10, making a £5 profit. Suppose there's a market for 100 widgets and both companies have equal market share. If one company lowers their price to £9, and profit to £4, they can increase their market share to 80% as most buyers will be price-savvy and choose to pay less. The following table represents the various outcomes.

| **Price War** | B cuts price to £9: profit £4 | B keeps price at £10, profit £5 |
|---|---|---|
| A cuts price | A's profit: £200<br>B's profit: £200 | A's profit: £320<br>B's profit: £100 |
| A maintains price | A's profit: £100<br>B's profit: £320 | A's profit: £250<br>B's profit: £250 |

Again looking from A's point of view, if B cuts the price the profit for A becomes either £200 or £100, so they should clearly cut the price too. If B keeps their prices the same, the profit for A becomes either £320 or £250, so again they should cut the prices. But once A cuts their price, the rational choice for B is to follow suit, and after it's all over, both companies again have equal market share, but now are only making £200 profit rather than £250. And this process could be repeated over and over until nobody is making any money at all. This, of course, is why collusion and price-fixing are such problems. Part of the solution is to convince buyers that your product is better and that it's worth paying more for than the competition's product.

One of the pioneers of game theory, David Blackwell, had an interesting comment about the prisoner's dilemma. Blackwell was an American mathematician whose key contributions were in the areas of game theory and statistics. He had originally planned to become a teacher, but then he took some courses in mathematics at university and realised this was the subject for him. He was awarded a Ph.D. in 1941 at only 22 years of age – only the seventh Ph.D. in mathematics awarded to an African American. He then moved to Princeton's Institute for Advanced Study for a year, even more impressive considering that at this time Princeton had never even had a Black student, never mind a Black faculty member. He was the first African-American elected member of the prestigious National Academy of Sciences. He chaired the University of California, Berkeley's Department of Statistics for 30 years (though the first time he applied for a post there he was turned down because the Head of Department's wife, who habitually hosted dinners for faculty members, refused to entertain the idea of having a person of colour in her house). He published more than 80 academic papers, and was highly influential in mathematics, with 50 Ph.D. students, well-regarded text books, and a reputation as an excellent teacher. In fact, he's one of the very few mathematicians that has made it into the pages of fiction – he's mentioned several times as an important mathematician and role model by the mathematician Odenigbo in Chimamande Ngozi Adichie's wonderful novel *Half of a Yellow Sun* (you can read about this and other mathematicians in fiction in my book *Once Upon a Prime*, by the way).

It was while working at the Rand corporation over the summers of 1948-50 that Blackwell's interest in games and strategies was developed. Unfortunately such ideas had many applications in that Cold War environment. Talking about the prisoner's dilemma, where the best outcome can only be achieved through cooperation, Blackwell said "[t]he situation with the Soviet Union has had elements like this in it. To cooperate is to disarm and to double-cross is to re-arm with bigger and bigger weapons. That takes a lot of resources and we would both be better off disarming. But each is afraid that if he throws away his weapons, the other one will not and he will be at a great disadvantage."

# How to fight a duel

David Blackwell is known for his work on the problem of duels. “One day some of us were talking and this question arose: If two people were advancing on each other and each one has a gun with one bullet, when would you shoot? If you miss, you're required to continue advancing. That's what gives it dramatic interest. If you fire too early your accuracy is less and there's a greater chance of missing.”

One version of the basic duel problem is this. Two duellists face each other 20 paces apart. They start walking towards each other, pistols raised. They may each fire a single bullet at each other, choosing when to do this. If one duellist hits the other first, he is the winner. If both miss, or both hit but at the same time, it's a draw. What is the best strategy?

As an example of the analysis, we will make some assumptions to simplify the problem. First, we assume that the duellists are equally skilled, so that from the same distance they have the same probability of hitting each other. Next, we move from continuous to discrete. That is, although in theory the duellists could fire at any of the infinitely many points along the path back to each other, we will only allow them actually to fire after a whole number of paces, between 1 and 10. We will also assume that accuracy is inversely proportional to distance apart. Therefore, we will say that if they fire after walking $p$ paces, then the chance of hitting their opponent is $\frac{p}{10}$. (Once each duellist has walked 10 paces, then the distance between them is zero, so it's impossible to miss.) Finally, if one duellist fires and misses before the other shoots, then the remaining duellist can simply walk up to his opponent and fire, guaranteeing that he will win.

To compare different strategies against one another, we use the idea of expected payoff. Let’s say we assign a payoff for each player of $+1$ for a win, $-1$ for a loss, and 0 if it’s a draw. This is a zero-sum game in the sense that if Rose wins, Colin loses, and vice versa, so the overall outcome is zero. We can work out the expected payoff for her, for each strategy the players might choose.

Suppose Rose decides to fire after 4 paces. If Colin fires after 1 pace, then Colin has fired first, and his chance of winning is $\frac{1}{10}$. So the expected payoff is $(-1)\times\frac{1}{10}+1\times\frac{9}{10}=\frac{8}{10}=0.8$. For 2 paces it’s 0.6, for three paces 0.4. If Colin fires after 4 paces, then they could both hit or both miss – both these have a value of 0 so we don’t need to worry about those in the calculation, so we only need to worry about the case where Colin hits and Rose misses, or vice versa. This gives a pay-off of $(-1)\times\frac{4}{10}\times\frac{6}{10}+1\times\frac{6}{10}\times\frac{4}{10}=0$. If Colin fires after five or more paces, then we know Rose has already fired. With probability $\frac{4}{10}$ she fires and hits, so wins. If she fires and misses, then Colin will definitely win. So the pay-off is $1\times\frac{4}{10}+(-1)\times\frac{6}{10}$. We can put this information in a table:

| Colin→ / ↓ Rose | 1 pace | 2 paces | 3 paces | 4 paces | 5 paces | 6 paces | 7 paces | 8 paces | 9 paces | 10 paces |
|---|---|---|---|---|---|---|---|---|---|---|
| 4 paces | 0.8 | 0.6 | 0.4 | 0 | $-0.2$ | $-0.2$ | $-0.2$ | $-0.2$ | $-0.2$ | $-0.2$ |

How do we determine the best-case scenario here? If we assign equal probability to each of Colin’s possible strategies, we can calculate an expected payoff by just averaging the possible payoffs: we get $\frac{1}{10}\left(0.8+0.6+0.4+0+6(-0.2)\right)=0.4$. So this is a positive payoff. But is it the best strategy? There are still six cases here where a loss is more likely than a win for Rose. One approach is called the maximin, or the pessimist’s philosophy. You look at the worst possible outcome for each choice of strategy, and you pick the outcome with the best “worst case scenario”. For the 4-pace strategy, the worst case is $-0.2$.

We can draw up the full table of possibilities (given over the page), and here I have just written the payoff for Rose in each entry. If her payoff is $p$, then Colin’s is $-p$, as it’s a zero-sum game. The column headings are the number of paces taken by Colin, the row headings are the number of paces taken by Rose.

One approach Rose could take is to find the average payoff in each case, but that would require an assumption about the probability of Colin applying each of his possible strategies. Another tactic would be to try and maximise the chance of winning amongst those worst case scenario minimum values (called a maximin strategy). We can see that if she does this she should fire after either 5 or 6 paces. To choose between R5 and R6, we can notice that R6 is “dominant” in the sense that whatever the relative probabilities of Colin’s possible strategies, R6 is at least as good if not better than R5, for Rose, in each case. So the best choice is R6. If Colin thinks in the same way, he will choose the “smallest maximum” because the higher the payoff, the better for Rose. So he will apply a minimax strategy and choose to fire after six paces as well. If either switches to a different strategy, their predicted payoff does not improve.

| | C1 | C2 | C3 | C4 | C5 | C6 | C7 | C8 | C9 | C10 | Worst |
|---|---|---|---|---|---|---|---|---|---|---|---|
| R1 | 0 | −0.8 | −0.8 | −0.8 | −0.8 | −0.8 | −0.8 | −0.8 | −0.8 | −0.8 | −0.8 |
| R2 | 0.8 | 0 | −0.6 | −0.6 | −0.6 | −0.6 | −0.6 | −0.6 | −0.6 | −0.6 | −0.6 |
| R3 | 0.8 | 0.6 | 0 | −0.4 | −0.4 | −0.4 | −0.4 | −0.4 | −0.4 | −0.4 | −0.4 |
| R4 | 0.8 | 0.6 | 0.4 | 0 | −0.2 | −0.2 | −0.2 | −0.2 | −0.2 | −0.2 | −0.2 |
| R5 | 0.8 | 0.6 | 0.4 | 0.2 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| R6 | 0.8 | 0.6 | 0.4 | 0.2 | 0 | 0 | 0.2 | 0.2 | 0.2 | 0.2 | 0 |
| R7 | 0.8 | 0.6 | 0.4 | 0.2 | 0 | −0.2 | 0 | 0.4 | 0.4 | 0.4 | −0.2 |
| R8 | 0.8 | 0.6 | 0.4 | 0.2 | 0 | −0.2 | −0.4 | 0 | 0.6 | 0.6 | −0.4 |
| R9 | 0.8 | 0.6 | 0.4 | 0.2 | 0 | −0.2 | −0.4 | −0.6 | 0 | 0.8 | −0.6 |
| R10 | 0.8 | 0.6 | 0.4 | 0.2 | 0 | −0.2 | −0.4 | −0.6 | −0.8 | 0 | −0.8 |

There are many variants of the dueling idea. You could make it continuous, so that players could fire at any point. You could change the level of skill of the players – if one is a better shot, how does it change things? And there’s another variant too. As David Blackwell recounted “It took us about a day to develop the theory of that duel…Then I got the idea of making each gun silent. With the guns silent, if you fire, the other fellow doesn't know, unless he's been hit. He doesn't know whether you fired and missed or whether you still have the bullet. That turned out to be a very interesting problem mathematically.” In this case there is an obvious application in the situation of warfare where you don’t necessarily know the other side has launched their missiles until they hit you.

## Finding an equilibrium

The duel is an example of a two-person zero-sum game, and one of the first major results in game theory is about such games. It was a theorem by John von Neumann, published in the paper *On the Theory of Games of Strategy* in 1928. I’d like to give a simpler example to illustrate it, and it’s to do with football (by which I mean soccer, for any American friends reading this).

In football, sometimes there are penalties. Here, the kicker can choose whether to aim left, right, or straight in the centre. The goalkeeper has the same choice to make. It’s a zero-sum game because if the ball goes in the back of the net the kicker “wins”, otherwise the keeper “wins”. Since 2009, according to Instat, https://instatsport.com/football/article/penalty_research, almost 100,000 penalty shots have been taken on football pitches around the globe. 75.49% of those resulted in goals, 17.57% were saved by goalkeepers, 4.07% went wide and 2.87% hit posts or crossbars. A higher proportion of shots aimed left or right are converted than shots aimed at the centre. But, of course, it’s useless to have a “pure” strategy of, say, always kicking right, because then everyone will get wise to it and the keepers will always move right. So, both the kicker and the keeper need to use some randomness – what’s known as “mixed” strategies.

Just as an example, you might observe some matches and find the following outcomes. The entries are probabilities. For example the number in the Keeper C, Kicker R entry means that if the kicker aims right and the keeper stays in the centre, then the probability of scoring is 0.8 (ie, 80%).

| | Keeper to (Kicker’s) Right | Keeper Centre | Keeper Left |
|---|---|---|---|
| Kicker aims Right | 0.5 | 0.8 | 0.9 |
| Kicker aims Centre | 0.8 | 0.3 | 0.8 |
| Kicker aims Left | 0.9 | 0.8 | 0.5 |

A given strategy will assign different probabilities to the different choices. Let’s suppose both players decide to choose a direction randomly with probability 1/3. What’s the expected average score? It will be $\frac{1}{9} \times 0.5 +$

$\frac{1}{9} \times 0.8 + \cdots + \frac{1}{9} \times 0.5 = 0.7$. On the other hand if the kicker switches to a strategy of going left half the time and right half the time (and never staying put), the expected score will be $\frac{1}{6} \times (0.5 + 0.8 + 0.9 + 0.9 + 0.8 + 0.5) = 0.73$. So this is better for the kicker, but the keeper will soon adapt their strategy to compensate. What von Neumann's Minimax Theorem says is that there is always an equilibrium strategy that's both a minimax strategy for one player and a maximin strategy for the other – and equilibrium here means that if either player unilaterally changes their strategy, the outcome for them becomes worse. So it's at least a local optimum. The equilibrium can be found using a technique called linear programming. A mathematician called Ferenc Forgó[1] analysed this situation and found there is an equilibrium if both players move right 42% of the time, left 42% of the time, and stay central 16% of the time. This gives an expected score of 0.72, in other words a goal is scored 72% of the time.

So far we've talked about two-person zero-sum games. And now John Nash enters the scene. He was born in 1928, the year John von Neumann's Minimax Theorem paper came out, and died in 2015. He is well known in popular culture because of the 1998 book A beautiful mind, by Sylvia Nassar, which was made into a 2001 film of the same name starring Russell Crowe, who I have to say doesn't look very much like John Nash, but never mind! The book and film portray Nash's struggles with mental illness. He was an undoubted genius – his 1950 PhD thesis was just 28 pages long and introduced groundbreaking ideas in game theory. The most famous of these is what is called the Nash equilibrium. He showed in a 2-page paper in the same year (Equilibrium points in N-person games) that the ideas of von Neumann's Minimax Theorem can be extended to non-zero-sum games with any number of players. A Nash equilibrium is a situation where no one player can change their strategy unilaterally without worsening their outcome. The prisoner's dilemma is a very simple example. If both players talk, then they both get 3 years in prison. If just one of them instead stays silent, then that player would instead get 5 years in prison. So the "both talk" outcome is a Nash equilibrium. The "best coffee in town" problem can be solved using these techniques. Say we have 400 locals and 200 tourists per day; the tourists choose randomly, so 100 will visit each café whatever the price, and the locals choose whichever café's prices are lower. Suppose each café can either charge £3 or £4 for a coffee. We can see that in this case, it's not a zero sum game, but there will be an equilibrium at £3 because if either café moves from this on its own, their revenue goes down from £900 to £400. If both cafés start at £4, then either café can increase profits by dropping the price to £3, so that's not an equilibrium. Nash's ideas have since been extended. For instance John Harsanyi looked at cases where some players have private information (like in an auction where bidders know the maximum they can afford, but not this information for other bidders). Rheinhardt Selten considerd what happens where some players can observe the actions of others before they make their own decisions. These three shared the 1994 Nobel Prize in Economics for their work.

Speaking of Nobel Prizes, the 2020 Nobel Prize for Economics was won by Paul Milgrom and Robert Wilson for their work on auction theory. So let's now talk about auctions.

# Auctions

There are several kinds of auctions.

- Open "English" auction (ascending bid). Starting from a minimum bid, the bid is increased until no further bids are made. The final bidder wins and the price paid is that final bid.
- Open "Dutch" auction (descending bid). Starting from a very high price, higher than anyone will actually pay, the auctioneer lowers the price until there is a bidder prepared to buy, at which point they win the auction at that price.
- Sealed bid auction: Everyone bids in secret; the highest bid wins. There are two variants: in the first-price variant the amount paid is the amount of the highest bid. In the second-price variant the amount paid is the amount of the second highest bid.

The duel problem is reminiscent of a Dutch-style auction, when a high price is initially set and then is brought down in increments until someone bids. A variant of this is often used on shopping channels – usually there is (claimed to be) a limited number of items on sale for one hour, and the price ticks down through the hour until they are all sold. The longer you wait, the better the bargain, but you risk losing out.

A lot of work has been done on auctions, and there are some interesting paradoxes. One of the problems here is that it's often hard to know the true value of an item. Different buyers may well have different

[1] See his paper at https://www.jstor.org/stable/10.2307/90002528

valuations -dealers and collectors for example, or perhaps those who have the expertise to mend a broken item, versus those who do not – that is, each buyer has a "private value" that they are assigning to the item, and they won't bid above that value. Others consider the case where the item objectively has a value (the "common value"), for example mining rights, but buyers will all have different estimates of that value (from different geological surveys, for instance). This latter scenario leads to what's known as the "winner's curse". By definition if you place the highest bid, then everyone else felt the item was not worth that much. The inescapable conclusion is that you overpaid.

It's in the design of auctions that game theory has the most to say. For instance, many auctions have a reserve price, where the item is not sold if that price is not reached. What is the best reserve price to set, in terms of the revenue you are going to get? If you set it too high, you risk coming away with nothing. Set it too low, and you may regret selling at all. We can reduce the problem to the case of just a single bidder (or think of it as the bidder prepared to pay the most out of a group of bidders). We don't know what that maximum bid will be, so we assume it will be a random point from between 0 and let's say £100 (we could set this based on prices achieved for similar items). Suppose our reserve price is R and the bid is $x$. If $0 \leq x < R$, we get nothing. If the maximum bid $x$ is greater than or equal to $R$, which happens for all values of $x$ between $R$ and 100, the revenue is $R$ (they would be prepared to pay more but obviously will offer the minimum amount needed to secure the item, so they will bid and pay $R$). So the expected average revenue is the area under the graph of bid against revenue, divided by 100. We get $\frac{1}{100}(R \times 0 + (100 - R) \times R) = \frac{1}{100}(100R - R^2)$. We can plot this graph to see how it varies with our choice of $R$, and we can see that it's an upturned parabola with the maximum value attained at $R = 50$. (You could also determine this using calculus or completing the square to see that $\frac{1}{100}(100R - R^2) = \frac{1}{100}(2500 - (R - 50)^2)$.

Actually, a similar calculation tells you what to bid in some situations, like a first-price sealed bid auction. Suppose you are a furniture dealer. You see a lamp, say, that you know you can refurbish and sell for £100. In a first-price sealed bid auction, you shouldn't bid £100 as then you make no profit. Assuming you have no insight into what other bidders may do, you should bid half what you think you can ultimately get for the lamp. In this case you should bid £50. This balances out the risk of loss of small profits when you don't win, with the chance of bigger profits when you do win.

A second-price sealed bid auction requires a different strategy, because the price the winner pays is actually the amount of the second-highest bid. Sealed bid auctions were first studied formally by William Vickrey in the 1960s, and second-price sealed bid auctions are also known as Vickrey auctions. What's the best bid to make? Let's say you value the item at £100. If you bid less, say £90, then someone might sneak in with a bid of £99 and you will regret your decision. If you bid more, say £110, then you might win the auction but have to pay a second-place bid amount of £109, so you will have overpaid, according to your valuation. The only regret-free choice is to bid precisely what you think it is worth. Of course the issue there is deciding accurately what you think an item is worth, which is rather difficult.

It's a rather curious fact that, due to something called the revenue equivalence theorem, the theoretical expected revenue is the same for all the auctions so far. But there are confounding factors in real life, for example the result ceases to follow if the buyers have budget constraints or are risk averse, or if the bids made are not statistically independent of each other (which is quite likely in real life). That expected winning price turns out to be half of the perceived value the winning bidder ascribes to the item. For English auctions, where it's open bidding and the highest bid wins, and is the price paid, let's think about the case of two bidders, Colin and Rose. They both will bid up to their valuation of the item. Suppose Rose wins. Then Colin's value must have been lower, or he would not have stopped bidding. If her value was £100, then Colin's value must be somewhere between 0 and £99.99. Unless we have further information the expected amount is just the average of all values in that range, which is the midpoint of £50. The point at which bidding stops is when Colin gives up, which is when bidding reaches his value (plus some small increment suppose). Therefore the expected selling price is £50, half Rose's valuation.

There are much more sophisticated game-theoretic calculations that have had significant real-world applications. One example happened in 2000 in the UK when the government auctioned off 3G airspace. Ken Binmore and Paul Klemperer led the team designing the auction, which raised 23 billion pounds – much higher than previous, similar auctions elsewhere. This result is great for the governments, and so hopefully the citizens, of the countries in question, though less great for the telecoms companies who overvalued what they were bidding for and ended up losing money.

## When to buy a house

There's a final question I want to tell you about, and it's the problem of renting or buying a house. If you are looking for a new place, you will view several properties. Certainly in the current rental market nothing stays available for long. So quite often if you don't put in an offer straightaway you'll lose the chance. It'd be silly to take the very first house you see as maybe there are loads of better ones out there. But when do you stop looking? If the house we just looked at is worse than one we've seen earlier, we won't really want to choose it, because we know for sure there are, or at least have been, better options. So our strategy ought to at least require that we only choose the current option if it's the best so far. If it *is* the best so far, then that gives us some hope, but there is a chance we might be about to see something even better, so we don't want to stop too soon.

We need some assumptions. We need to make an estimate of the total number $n$ of houses we are prepared to view. We assume that in any collection the houses can be ranked best to worst with no ties, and that the houses are randomly ordered in terms of when we view them, so that there's a $1/n$ chance that the first house we see is best overall, a $1/n$ chance that the second house is best, and so on. Finally, our strategy will be to view a certain number of houses to get a sample, and then take the first house after this that's the best so far. If we never see a better house then we either don't buy at all, or we end up with the last house we see. So the strategy will fail, for example, if the best house overall was in the sample set.

If we try this with three houses, then our sample size could be 0, 1, or 2 (but not 3 as we are rejecting these sample houses).

- If $s = 0$, then House 1 is the best, by virtue of being the only, house we've seen. So we take it. There's a 1/3 chance it's the best overall. So the success rate of this strategy is 1/3.
- If $s = 1$, we reject House 1, then adopt the "choose the next house that's better than any previous houses" strategy.
  - If House 1 is the best (a 1/3 chance), then we reject House 2 and House 3, so the strategy fails.
  - If House 2 is the best (a 1/3 chance), then we choose House 2 as it's the best so far at that point, so the strategy succeeds.
  - If House 3 is the best then there are two cases. If the ranking is House 1 worst, House 2 better, House 3 best (1/6 chance), then we choose House 2, so the strategy fails. If it's House 2 worst, House 1 better, House 3 best (1/6 chance) then we choose House 3, so the strategy succeeds.

  Therefore, overall, there is a $\frac{1}{3} + \frac{1}{6} = \frac{1}{2}$, or a 50% chance of success by setting $s = 1$.
- If $s = 2$, then we reject the first two houses automatically, so the strategy only succeeds if House 3 is the best, again a 1/3 success rate.

The best strategy then is $s = 1$: reject the first 1 of the 3 houses, or 33%, and pick the best so far after that. For higher numbers of houses, you can test the different strategies and look for a pattern.

| Number of houses | best sample size | best sample size % | best sample size fraction |
|---|---|---|---|
| 3 | 1 | 33.3% | $^1/_3$ |
| 4 | 1 | 25% | $^1/_4$ |
| 5 | 2 | 40% | $^1/_{2.5}$ |
| 6 | 2 | 33.3% | $^1/_3$ |
| 7 | 2 | 28.6% | $^1/_{3.5}$ |
| 8 | 3 | 37.5% | $^1/_{2.7}$ |
| 9 | 3 | 33.3% | $^1/_3$ |
| 10 | 3 | 30.0% | $^1/_{3.3}$ |
| Large | | 36.79% | $^1/_{2.718}$ |

It turns out that, for large $n$, the best sample size is around 36.79% of the total expected number, or as a proportion, around 1 in 2.718. The exact number can be found using a bit of calculus (there's a link in the further reading that has some details of this), and is $1/e$, the famous mathematical constant. The success rate is also $1/e$, which is pretty good given we have no idea where our best house will be in the distribution. There's another potential application of this idea too, and that's to dating. How many people should you date before settling down? To answer this, you have to estimate the total number of potential partners – maybe on average people are dating in their 20s and plan to marry or settle down at some point in that time, so that's up to ten years, and perhaps we can allow 2 relationships a year, so we have a pool of 20 candidates; from this we should not marry any of the first seven, but from number 8 onwards we marry anyone who is the best so far.

In our journey through the mathematics of choice and games, we've seen that these ideas have a huge number of applications both in economics and in life in general. I hope you've enjoyed it.



## Further Reading

- Ken Binmore, who was instrumental in the 3G auctions we discussed, has written an excellent non-technical book *Game Theory: A Very Short Introduction* (Oxford University Press 2007).
- For a more academic introductory text, *Game Theory Through Examples*, by Erich Prisner (Cambridge University Press 2014) is recommended for self-study. It's very approachable and, as the title suggests, is full of examples.
- For a biography of John von Neumann, try Ananyo Bhattacharya's book *The Man from the Future: The Visionary Life of John von Neumann* (Penguin, 2021).
- The classic biography of John Nash is of course *A Beautiful Mind*, by Sylvia Nasar (Simon & Schuster, 1998). You could also watch the 2001 film starring Russell Crowe, which is very entertaining.
- While there is a certain amount of information online about David Blackwell, I've not been able to find a book about him, which is a shame. The Wikipedia article links to various obituaries and magazine articles, and these are a good place to start. https://en.wikipedia.org/wiki/David_Blackwell
- *Half of a Yellow Sun*, by Chimamande Ngozi Adichie, tells the story of the Biafran war. It was published in 2006 by 4th Estate.
- My book *Once Upon a Prime: The Wondrous Connections between Mathematics and Literature* is published in 2023 by Mudlark in the UK and Flatiron in the US.
- There's a nice write-up of the mathematical derivation of the $\frac{1}{e}$ solution to the house-buying problem, in its guise as the "Secretary problem", at https://datagenetics.com/blog/december32012/index.html
- I drew the graphs illustrating the reserve price problem and the house-buying problem using the free online graphing tool Desmos. https://www.desmos.com/calculator

## Image Credits

Images are either drawn by me or are publicly available/free use; the photograph of John von Neumann is from the Los Alamos National Laboratory and is free to use subject to the requirement to include this text:

"Unless otherwise indicated, this information has been authored by an employee or employees of the Los Alamos National Security, LLC (LANS), operator of the Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396 with the U.S. Department of Energy. The U.S. Government has rights to use, reproduce, and distribute this information. The public may copy and use this information without charge, provided that this Notice and any statement of authorship are reproduced on all copies. Neither the Government nor LANS makes any warranty, express or implied, or assumes any liability or responsibility for the use of this information."