DOI: 10.13140/RG.2.2.31675.98089

The Deer Hunter: A lesson in the basics of risk and probability assessment

Norman Fenton[1]

Queen Mary University London and Agena Ltd

25 April 2020

I was recently watching a re-run of the classic 1978 Michael Cimino film "The Deer Hunter". It contains one of the most iconic scenes in cinema history involving a 'game' of Russian roulette forcibly played by two American soldiers held captive in Vietnam (https://youtu.be/4wv2K3J__X0). Although I have seen the film several times, this scene never seems to lose its impact. As I am currently teaching a new course on Risk Assessment and Decision Making, it also occurred to me that the scene provides a rich source of examples to illustrate core concepts of probability and risk including: probability and odds, basic probability axioms, conditional probability, risk and utility, absolute versus relative risk, event trees, and Bayesian networks.

The context

Life-long friends Mike (played by Robert DeNiro) and Nick (played by Christopher Walken) are captured American soldiers in Vietnam forced to play Russian roulette by their Viet Cong guards. The 'normal' game - involving a single bullet and two or more participants – usually ends when a participant is killed. Mike and Nick have previously already played and survived. They are now forced to play again - this time against each other.

Realising that one of them is very likely to die, Mike proposes an alternative version of the 'game' – involving three bullets - which he believes will enable them both to escape by killing the guards.

As the iconic clip https://youtu.be/4wv2K3J__X0 shows, Mike's audacious plan works.

Figure 1 Still from the scene

The plan raises some very interesting questions about the 'risk' Mike was taking and the probability of success. In fact, many basic concepts in risk and probability are easily illustrated by comparing the 'risks' involved in the two different games.

[1] n.fenton@qmul.ac.uk, Twitter: @ProfNFenton

The Two Games

Figure 2 The 'normal' game involves a single bullet in one of the six chamber of the gun barrel

Note that pulling the trigger moves the barrel on to the next chamber (in a less interesting variation of the game the gun barrel is spun after each time the trigger is pulled, meaning that each time we have the same probability 1/6 of firing the bullet)

Mutually exclusive outcomes or events

A set of events or outcomes are ***mutually exclusive*** if at most one can happen.

Probability of an event

For any event or outcome *X* we write the probability of *X* as Prob(*X*). So, if *X* is the event "roll a 4 on a 6-sided die"

Prob(*X*)=1/6

All probabilities range from 0 to 1 (although they are sometimes converted to percentages by multiplying by 100), with 0 representing certainty the event does not occur and 1 representing certainty the event does occur.

The 'normal game' involves a single bullet loaded into one of the six chambers of a gun barrel. The barrel is spun before the game starts to ensure the chamber with the bullet is randomly located.

The gun is then spun on the table to determine which player goes first (namely the player the gun points to). The player must pull the trigger with the gun pointed at their own head. If there is no bullet in the chamber when the trigger is pulled, the next player must point the gun at their own head and pull the trigger. The game continues until a bullet is fired (and in what follows we assume this always results in the death of the shooter). The only exception is when there is no bullet in the first five chambers triggered, i.e. when the only remaining chamber must contain the bullet. In this case the game ends with nobody shot.

Those watching the game gamble on the outcome either at the beginning of the game or after each time the trigger is pulled without a bullet being fired. The reason why the game stops if the only remaining chamber must contain a bullet is because there is no uncertainty of the outcome and so no interest from a gambling perspective.

The 'alternative game' proposed by Mike involves **three** bullets, rather than one, loaded (randomly) into the gun barrel. Otherwise the rules are the same (so again, in the unlikely event that the first three chambers are blank, the game ends with nobody shot).

The possible outcomes (and their probabilities) for the 'normal' game

For the normal one-bullet game played with two players there are three possible (*mutually exclusive*) outcomes:

	**A. Player 1 dies**	☹
	**B. Player 2 dies**	☹
 	**C. Neither player dies**	☺

To calculate the probabilities of these outcomes we must consider the different 'events' that can lead to them (an outcome is just a collection of events). In fact, there are six such mutually exclusive events:

E1. Player 1 dies on the 1st trigger press
E2. Player 2 dies on the 2nd trigger press
E3. Player 1 dies on the 3rd trigger press
E4. Player 2 dies on the 4th trigger press
E5. Player 1 dies on the 5th trigger press
E6. Neither player dies

Each of these events is completely determined by the location of the bullet in the barrel (where chamber 1 represents the start chamber, which moves to chamber 2 when the trigger is pressed etc):

Bullet location	**Event**	**Probability**
Chamber 1	E1. Player 1 dies on the 1st trigger	1/6
Chamber 2	E2. Player 2 dies on the 2nd trigger	1/6
Chamber 3	E3. Player 1 dies on the 3rd trigger	1/6
Chamber 4	E4. Player 2 dies on the 4th trigger	1/6
Chamber 5	E5. Player 1 dies on the 5th trigger	1/6
Chamber 6	E6. Neither player dies	1/6

Since the one bullet is randomly located in one of the 6 chambers, the probability it is in any specific chamber is simply 1/6. This means that, at the start of the game, each of the six events has the same probability 1/6.

Now outcome A (player 1 dies) is made up of the 3 mutually exclusive events E1, E3 and E5. Hence (see sidebar):

Prob(A) = Prob(E1 or E3 or E5)

= Prob(E1)+Prob(E3)+Prob(E5) = 1/6 + 1/6 +1/6 = 1/2

Similarly,

Prob(B) = Prob(E2 or E4) = Prob(E2) + Prob(E4) =1/6 + 1/6 = 1/3

Finally,

Prob(C) = Prob (E6) = 1/6

So, at the start of the game, the rational betting strategy (assuming 'fair odds' are offered – see sidebar) is to bet on Player 1 dying, since this is the most likely outcome. It is important to note, however, that the probability of all the events change (as explained in the sidebar 'conditional probabilities') as soon as we observe the outcome of E1 (and then again after further observed outcomes).

Of course, as Mike and Nick are very close friends, it is only outcome C that is a 'good outcome' for them. Because the probability of a "a player dies" is the negation of event C, it follows (see sidebar) that:

Prob("a player dies") = Prob(not C) = 1 - Prob(C) = 5/6

The probability of a set of mutually exclusive events

By an axiom of probability this is simply the sum of the individual probabilities. So if *X* and *Y* are mutually exclusive:

Prob(*X* or *Y*) = Prob(*X*)+Prob(*Y*)

Odds version of probability

If Prob(*X*)=1/3 then we also say the: "the odds are 2 to 1 against *X*"

In general, if Prob(*X*)= *a*/*b* then: the odds are "(*b*-*a*)/*a*" against *X*

Conditional and marginal probabilities.

Although the 'starting' probability of an event like E2 is 1/6, once we *know* whether the preceding event E1 is true or false the probability of E2 changes. For example, if E1 is false then the probability of E2 is 1/5. This called the ***conditional probability*** of E2 given not E1. We write this as

Prob (E2 | not E1) = 1/5

In contrast, Prob(E2), which we know is 1/6, is formally called the ***unconditional*** or ***marginal probability*** of E2.

Negation of an event or outcome

Another axiom of probability concerns the probability of the negation of an event or outcome *X* which is written as "not *X*". Specifically, the probability of the negation of an event *X* is one minus the probability of *X*, i.e.

Prob (not *X*) = 1 – Prob(*X*)

We could also have got to the same result by noting that the outcome "a players dies" is made up of the mutually exclusive events A and B whose probabilities 1/2 and 1/3 respectively also sum to 5/6.

We can think of the probability of C, i.e. 5/6, as providing a baseline for the ***risk*** to the players. But – as we shall see – risk should also involve the notion of utility of outcomes. While any outcome other than C clearly has catastrophic negative utility, even outcome C has negative utility since it is only likely to delay both their deaths, while resulting in further torture. This will become clearer when we consider the alternative 3-bullet game and its possible outcomes, because we need to understand why Mike felt it was worth 'taking the risk' of insisting on playing the alternative game.

Appendices 1 and 2 provide different ways of arriving at the above probability solutions for the one-bullet game (one uses the *event tree method* and the other use a *Bayesian network*).

The possible outcomes (and probabilities) for the 'alternative' 3-bullet game

We have the same three outcomes we had for the normal game. However, the 'good outcome' C has a very different meaning, namely the one that Mike's escape plan relies on. Specifically, this requires the following sequence of events all to be true[2]:

1. Gun is spun to point to Mike (so Mike is chosen to pull the first trigger)
2. No bullet in first chamber (So Mike triggers a blank)
3. No bullet in second chamber (so Nick triggers a blank)
4. There is a bullet in the third chamber that Mike uses to fire at the lead guard rather than at his own head.
5. Mike uses the remaining two bullets to kill two other guards, enabling him and Nick to take their weapons and escape (we assume the remaining two bullets do not have to be in the next two chambers, i.e. there could be a gap).

The simplest way of explaining the probability calculations in this case is to use the event tree shown in Figure 3 (an equivalent Bayesian network solution is provided in Appendix 3).

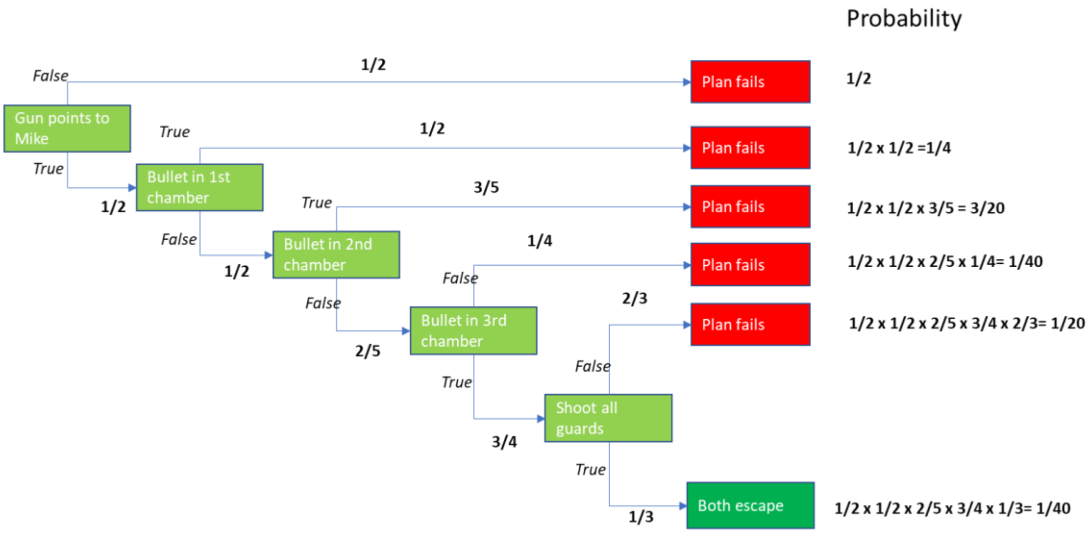

Figure 2 Event tree showing probabilities for the alternative 3-bullet game

[2] Note that, although theoretically both survive if the first 3 chambers are empty, Mike does not consider this to be a 'good outcome'.

Event Trees

An event tree shows the sequence of events leading to each different outcome of interest. Each branch corresponds to whether the event it is linked from is true or false. The branch is labelled with either true or false and its probability.

Computing the branch sequence probabilities

To compute the probability of any outcome such as any of the "plan fails" outcomes or the "Escape" outcome we multiply the branch probabilities leading to that outcome. This actually follows from another probability axiom which says that if an event X is comprised of a set of independent events X1, X2, .., Xn then the probability of X is the product of the individual event probabilities.

The probability for the event "Shoot all guards"

We can only use subjective probabilities for this. What we know from earlier in the film is that Mike is an expert shooter (as the name of the film indicates, he is a deer hunter) and Nick is also very competent. But there are several guards, so we assume a probability of 1/3 (and note again that this probability is conditioned on all the 'right' outcomes of the previous steps).

This tree is based around the sequence of 5 events required for the plan to succeed. If any one of the events is false then the Plan fails. Hence, we have 5 separate "Plan fails" outcomes (we can refer to these as F1, F2, F3, F4, F5 respectively). So, if the first event "Gun points to Mike" is false the plan immediately fails, i.e. we arrive at F1. The second event "Bullet in 1st chamber" is only relevant if the first event was true. If "Bullet in 1st chamber" is true then the plan fails, i.e. we arrive at event F2 etc.

Most of these probabilities are straightforward; for example, the probability there is a bullet in the first chamber is 1/2 because at the start of the game we know there are 3 bullets randomly located in 6 chambers. But it is important to note that all the probabilities are conditional probabilities. For example, the reason for the probability 3/5 in the "True" branch of "bullet in 2nd chamber" is because this event lies on a branch in which there was no bullet in the 1st chamber. So, let B1 be "bullet in 1st chamber" and B2 be "bullet in 2nd chamber". Then the "True" branch from B2 represents the conditional probability, i.e. Prob (B2 | not B1) = 3/5

As explained in the sidebar for the normal game, this is NOT the same as Prob(B2), as the unconditional/marginal probability of B2 is 1/2.

The only contentious branch probability is the last, i.e. "Shoot all guards". For reasons explained in the sidebar we assume this is 1/3.

As explained in the sidebar, to compute the outcome probabilities we multiply the probabilities on each branch of the sequence leading to that outcome. So, for example the "Escape" outcome is comprised of the following "True" events (each conditioned on the previous ones)

- Gun points to Mike, which has probability 1/2
- Not B1, which has conditional probability 1/2
- Not B2, which has conditional probability 2/5
- B3, which has conditional probability 3/4
- Shoot all guards, which has conditional probability 1/3

Hence, Prob(C) = Prob(escape) = 1/2 x 1/2 x 2/5 x 3/4 x 1/3 = 1/40

So, the probability that Mike's plan works is just 1/40, meaning there is a 39/40 probability that his plan fails.

Probability versus risk: Did Mike make a rational choice?

If we define risk purely by the probability of the 'bad' outcome (as many people do), then Mike's decision to choose the alternative game rather than the normal game does not seem to make sense. The probability of a good outcome has dropped from 1/6 (16.67%) to

1/40 (2.5%) while the probability of a bad outcome has increased from 5/6 (=0.8333) in the normal one-bullet game to 39/40 (=0.975) in Mike's alternative three-bullet game.

Relative versus absolute risk increase

If the probability of dying from lung cancer is 0.04 (i.e. 4%) for non-smokers and 0.06 (i.e. 6%) for smokers then the ***relative risk*** of dying from lung cancer is 50% more for smokers than for non-smokers. However, the probability of dying from lung cancer is still small for smokers. Hence, it is more meaningful to use the ***absolute risk*** measure which is simply the absolute increase in probability. In this case absolute risk increase for smoker compared to non-smokers is 0.02, i.e. 2%.

Expected utility

Suppose a decision has a 'good outcome' and a 'bad outcome'. Then the expected utility of the decision is defined as

Prob (good outcome) x Utility (good outcome)

+

Prob (bad outcome) x Utility (bad outcome)

These risk changes are typically measured either as change in **relative risk** or **absolute risk** (see sidebar):

- The **relative risk** increase, which is (0.975 – 0.8333)/0.8333) is 17%.
- The more meaningful **absolute risk** increase, which is (0.975 – 0.8333) is 14%

Mike's decision only starts to make sense when we incorporate the **utility** (or its negation **cost**) of the outcomes. Crucially, the utility of a good outcome differs greatly for the two games, while the cost of a bad outcome is essentially the same.

If the 'cost' of a lost life is, say $5 million, then the utility of a bad outcome in the normal game is -$5 million. The bad outcome in the alternative game is essentially the same (although there is a very small probability that both Mike and Nick could die).

However, the good outcome in the normal game still results in Mike and Nick imprisoned, tortured, and facing almost certain death in future games. So, it also has a negative utility, but instead of -$5 million, let us say -$4 million, allowing for an unlikely imminent rescue. In contrast, the good outcome for the alternative game results in freedom for both Mike and Nick. We assume this is worth $10 million.

When there are choices (such as between the normal game and the alternative game) the optimal decision choice is the one which maximizes the overall expected utility (see sidebar). For the normal game this computes as:

$$\left(\frac{1}{6} \times -4{,}000{,}000\right) + \left(\frac{5}{6} \times -5{,}000{,}000\right) = -666{,}666 - 4{,}1666{,}666$$
$$= -4{,}833{,}333$$

For the alternative game this computes as:

$$\left(\frac{1}{40} \times 10{,}000{,}000\right) + \left(\frac{39}{40} \times -5{,}000{,}000\right) = 250.000 - 4{,}875{,}000$$
$$= -4{,}625{,}000$$

So, the optimal choice was indeed the alternative game - despite its low probability of success – because it has a higher expected utility (i.e. lower expected cost). Despite the low probability of success for Mike's plan, he made a rational decision to attempt it.

Further Reading

Fenton, N. E., & Neil, M. (2018). Risk Assessment and Decision Analysis with Bayesian Networks (2nd ed.). CRC Press, Boca Raton.

Spiegelhalter, D. (2019). The Art of Statistics: Learning from Data. Pelican Books

Appendix 1: Event Tree representation of the normal game

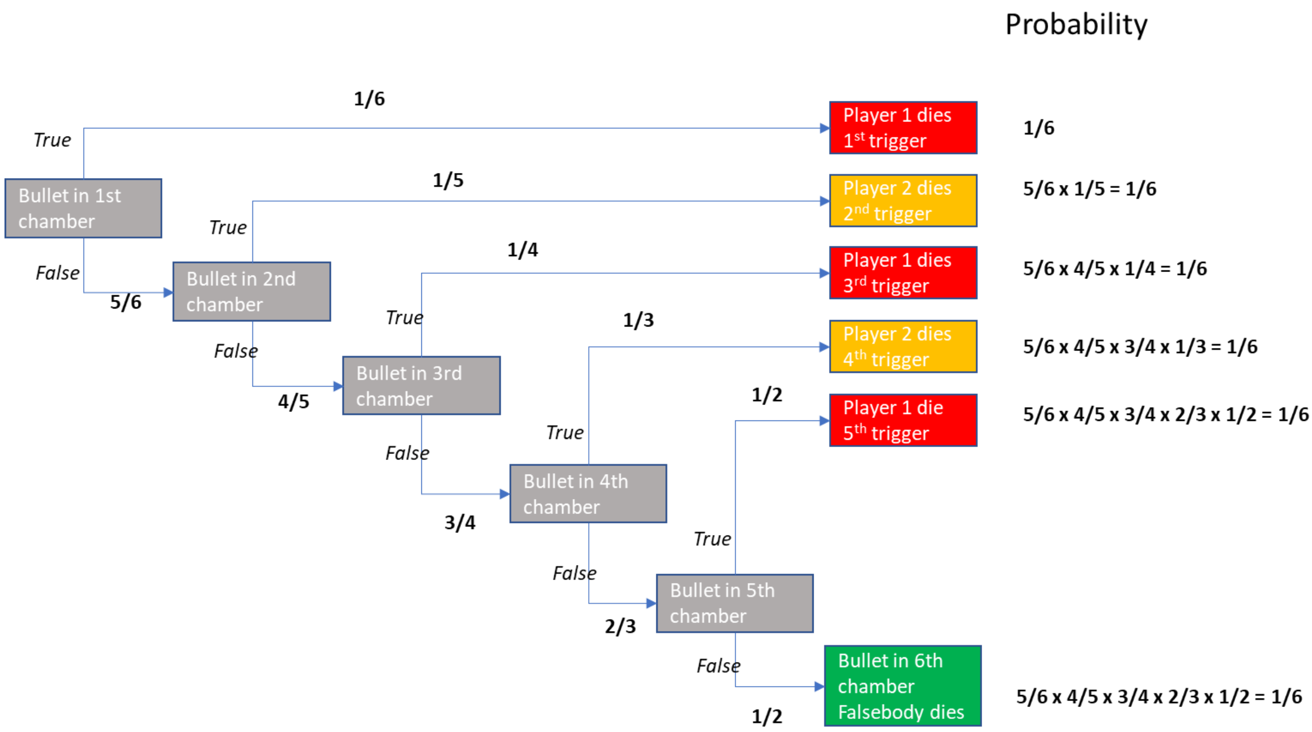

Figure 3 Event tree for 'normal' game

Appendix 2: Bayesian network representation of the normal game

The fully specified Bayesian Network (BN) is shown in Figure 5.

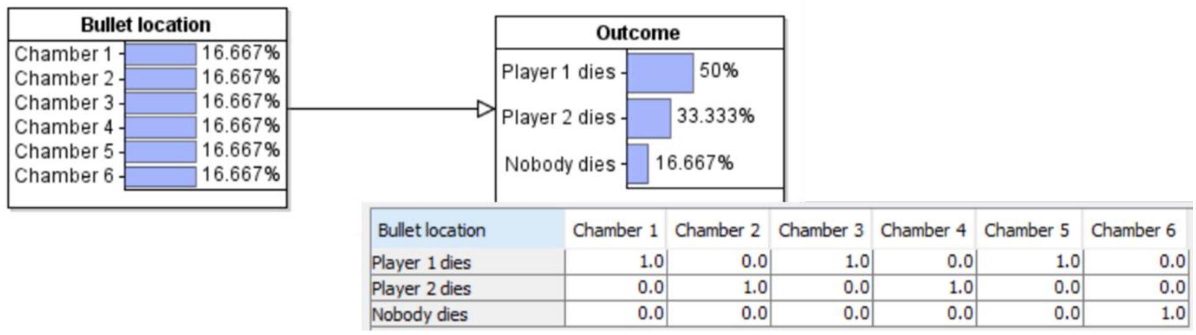

Bullet location	Chamber 1	Chamber 2	Chamber 3	Chamber 4	Chamber 5	Chamber 6
Player 1 dies	1.0	0.0	1.0	0.0	1.0	0.0
Player 2 dies	0.0	1.0	0.0	1.0	0.0	0.0
Nobody dies	0.0	0.0	0.0	0.0	0.0	1.0

It Figure 4 BN model

It is based on two variables – the *bullet location*, and the *Outcome*. The 'prior' probabilities for the 6 bullet locations are all 1/6. The conditional probability table for *Outcome* is shown and is self-explanatory. Once these conditional probabilities are entered the BN tool automatically calculates the marginal probabilities for the Outcome states. The BN solution is not only extremely compact and simple, but it enables us to perform different types of analysis. For example, we can automatically perform backward inference by entering an observation in the *Outcome* node as shown in the example in Figure 6.

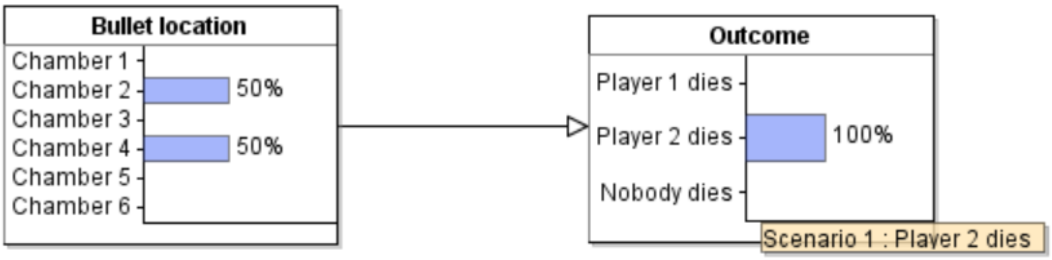

Figure 5 Backward inference. Computing the probability distribution for Bullet location given a particular outcome

Appendix 3: Bayesian network representation of the 3-bullet game