FRM Part I Exam

By AnalystPrep

Study Notes - Quantitative Analysis

Last Updated: Mar 7, 2026

Table of Contents

Reading 1: Fundamentals of Probability

After completing this reading, you should be able to:

- Describe an event and an event space.
- Describe independent events and mutually exclusive events.
- Explain the difference between independent events and conditionally independent events.
- Calculate the probability of an event for a discrete probability function.
- Define and calculate a conditional probability.
- Distinguish between conditional and unconditional probabilities.
- Explain and apply Bayes' rule.

Probability is the foundation of statistics, risk management, and econometrics. Probability quantifies the likelihood that some event will occur. For instance, we could be interested in the probability that there will be a defaulter in a prime mortgage facility.

Sample Space, Event Space, and Events

Sample Space (Ω)

A sample space is defined as a collection of all possible occurrences of an experiment. The outcomes are dependent on the problem being studied. For example, when modeling returns from a portfolio, the sample space is a set of real numbers. As another example, assume we want to model defaults in loan payment; we know that there can only be two outcomes: either the firm defaults or it doesn't. As such, the sample space is Ω = {Default, No Default}. To give yet another example, the sample space when a fair six-sided die is tossed is made of six different outcomes:

$\Omega = \{1, 2, 3, 4, 5, 6\}$

Events (ω)

An event is a set of outcomes (which may contain more than one element). For example, suppose we tossed a die. A "6" would constitute an event. If we toss two dice simultaneously, a {6, 2} would constitute an event. An event that contains only one outcome is termed an **elementary event.**

Event Space (*F*)

The event space refers to the set of all possible outcomes and combinations of outcomes. For example, consider a scenario where we toss two fair coins simultaneously. The following would constitute our event space:

{HH, HT, TH, TT}

Note: If the coins are fair, the probability of a head, P(H), equals the probability of a tail, P(T).

Probability

The probability of an event refers to the likelihood of that particular event occurring. For example, the probability of a Head when we toss a coin is 0.5, and so is the probability of a Tail.

According to frequentist interpretation, the term probability stands for the number of times an event occurs if a set of independent experiments is performed. But this is what we call the frequentist interpretation because it defines an event's probability as the limit of its relative frequency in many trials. It is just a conceptual explanation; in finance, we deal with actual, non-experimental events such as the return earned on a stock.

Independent and Mutually Exclusive Events

Mutually Exclusive Events

Two events, A and B, are said to be mutually exclusive if the occurrence of A rules out the occurrence of B, and vice versa. For example, a car cannot turn left and turn right at the same time.

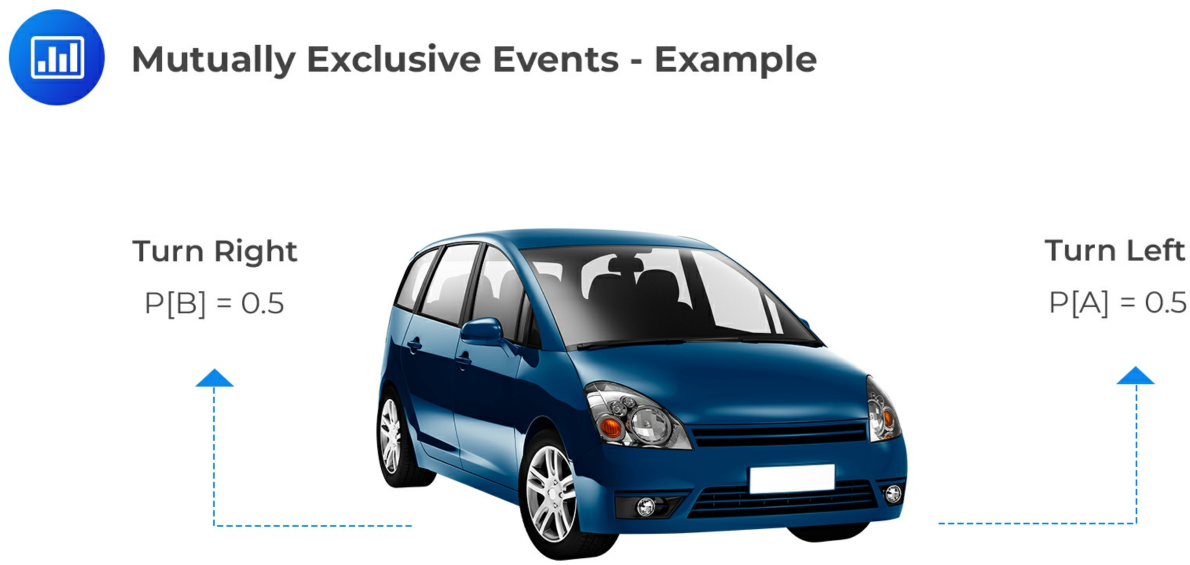

Mutually exclusive events are such that one event precludes the occurrence of all the other events. Thus, if you roll a dice and a 4 comes up, that particular event precludes all the other events, i.e., 1,2,3,5 and 6. In other words, rolling a 1 and a 5 are mutually exclusive events: they cannot occur simultaneously.

Furthermore, there is no way a single investment can have more than one arithmetic mean return. Thus, arithmetic returns of, say, 20% and 17% constitute mutually exclusive events.

Independent Events

Two events, A and B, are independent if the fact that A occurs does not affect the probability of B occurring. When two events are independent, this simply means that both events can happen at the same time. In other words, the probability of one event happening does not depend on whether the other event occurs or not. For example, we can define A as the likelihood that it rains on March 15 in New York and B as the probability that it rains in Frankfurt on March 15. In this instance, both events can happen simultaneously or not.

Another example would be defining event A as getting tails on the first coin toss and B on the second coin toss. The fact of landing on tails on the first toss will not affect the probability of getting tails on the second toss.

Intersection

The intersection of events say A and B is the set of outcomes occurring both in A **and** B. It is denoted as P(A∩B). Using the Venn diagram, this is represented as:

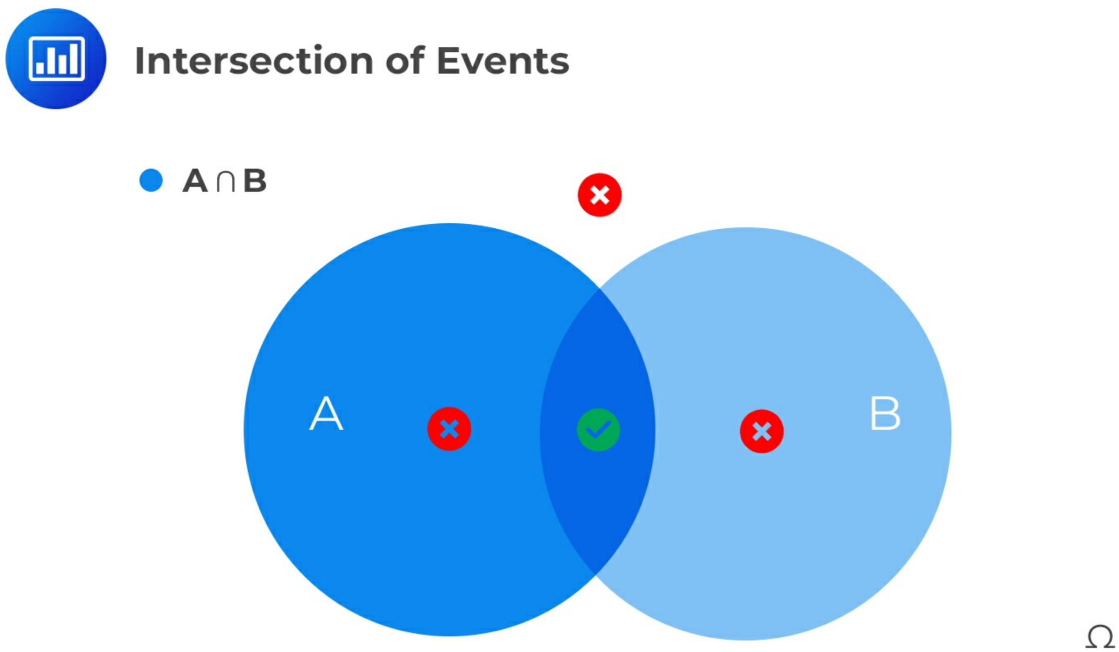

For independent events,

$$P(A \cap B) = P(A \text{ and } B) = P(A) \times P(B)$$

Independence can be extended to n independent events: Let $A_1, A_2, \ldots, A_n$ be independent events then:

$$P(A_1 \cap A_2 \cap \ldots \cap A_n) = P(A_1) \times P(A_2) \times \ldots \times P(A_n)$$

For mutually exclusive events,

$$P(A \cap B) = P(A \text{ and } B) = 0$$

This is because A's occurrence rules out B's occurrence. Remember that a car cannot turn left and turn right at the same time!

Union

The union of events, say, A and B, is the set of outcomes occurring in at least one of the two sets – A **or** B. It is denoted as P(A∪B). Using the Venn diagram, this is represented as:

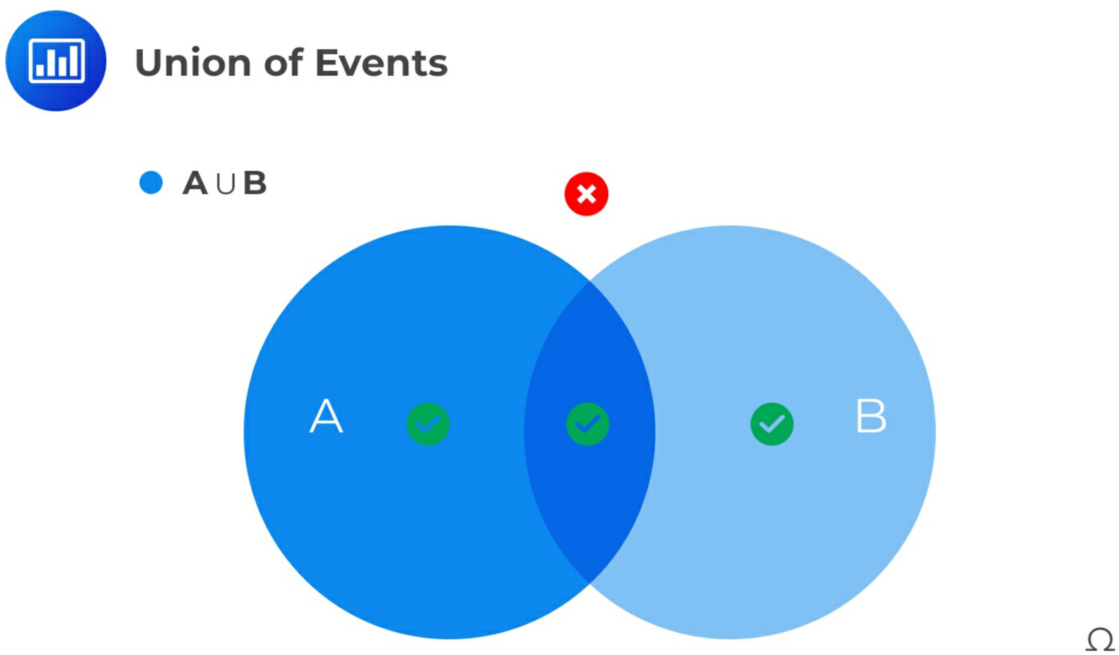

To determine the likelihood of any two **mutually exclusive events** occurring, we sum up their individual probabilities. The following is the statistical notation:

$$P(A \cup B) = P(A \text{ or } B) = P(A) + P(B)$$

Given two events A and B, that are not mutually exclusive (**independent events**), the probability that **at least one** of the events will occur is given by:

$$P(A \cup B) = P(A \text{ or } B) = P(A) + P(B) - P(A \cap B)$$

The Complement of a Set

Another important concept under probability is the **complement of a set** denoted by **Ac** (where **A** can be any other event) which is the set of outcomes that are not in A. For example, consider the following Venn diagram:

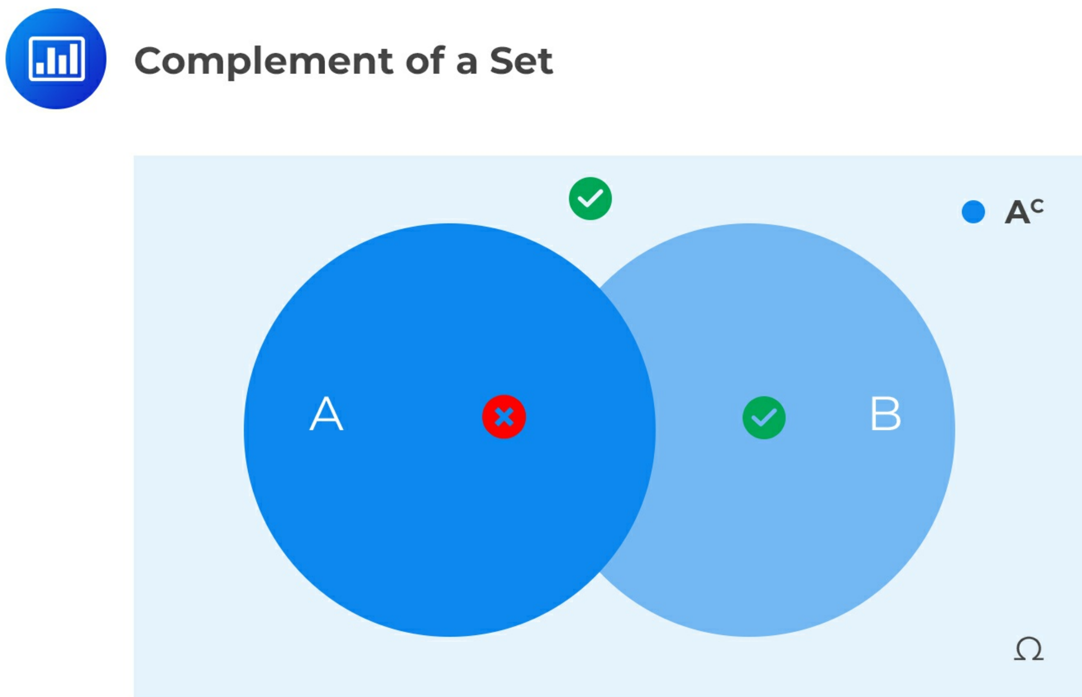

This is the first axiom of probability, and it implies that:

$$P(A \cup A^c) = P(A) + P(A^c) = 1$$

Conditional Probability

Until now, we've only looked at unconditional probabilities. An **unconditional probability** (also known as a marginal probability) is simply the probability that an event occurs without considering any other preceding events. In other words, unconditional probabilities are **not** conditioned on the occurrence of any other events; they are 'stand-alone' events.

Conditional probability is the probability of one event occurring with some relationship to one or more other events. Our interest lies in the probability of an event 'A' **given** that another event 'B **'has already occurred.** Here’s what you should ask yourself:

"What is the probability of one event occurring **if** another event has already taken place?" We pronounce P(A | B) as "the probability of A given B.," and it is given by:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

The bar sandwiched between A and B simply indicates "given."

Bayes' Theorem

Bayes' theorem describes the probability of an event based on prior knowledge of conditions that might be related to the event. Assuming that we have two random variables, A and B, then according to Bayes' theorem:

$$P(A|B) = \frac{P(B|A) \times P(A)}{P(B)}$$

Applying Bayes' Theorem

Supposing that we are issued with two bonds, A and B. Each bond has a default probability of 10% over the following year. We are also told that there is a 6% chance that both the bonds will default, an 86% chance that none of them will default, and a 14% chance that either of the bonds will default. All of this information can be summarized in a probability matrix.

Often, there is a high correlation between bond defaults. This can be attributed to the sensitivity displayed by bond issuers when dealing with broad economic issues. The 6% probability of both the bonds defaulting is higher than the 1% probability of default had the default events been independent, (I., P(A) × P(B))

The features of the probability matrix can also be expressed in terms of conditional probabilities. For example, the likelihood that bond A will default given that B has defaulted is computed as:

$$P(A|B) = \frac{P[A \cap B]}{P[B]} = \frac{6\%}{10\%} = 60\%$$

This means that in 60% of the scenarios in which bond B will default, bond A will also default.

The above equation is often written as:

$$P[A \cap B] = P(A|B) \times P[B] \qquad \text{I}$$

Also:

$$P[A \cap B] = P(B|A) \times P[A] \qquad \text{II}$$

Both the right-hand sides of equations I and II are combined and rearranged to give the Bayes' theorem:

$$P(B|A) \times P[A] = P(A|B) \times P(B)$$

$$\Rightarrow P(A|B) = \frac{P(B|A) \times P[A]}{P[B]}$$

When presented with new data, Bayes' theorem can be applied to update beliefs. To understand how the theorem can provide a framework for how exactly the new beliefs should be, consider the following scenario:

Example: Applying Baye's Theorem

Based on an examination of historical data, it's been determined that all fund managers at a certain Fund fall into one of two groups: Stars and Non-Stars. Stars are the best managers. The probability that a Star will beat the market in any given year is 75%. Other managers are just as likely to beat the market as they are to underperform it [i.e., Non-Stars have 50/50 odds of beating the market. For both types of managers, the probability of beating the market is independent from one year to the next. Stars are rare. Of a given pool of managers, only 16% turn out to be Stars.

A new manager was added to the portfolio of funds three years ago. Since then, the new manager has beaten the market every year. What was the probability that the manager was a star when the manager was first added to the portfolio? What is the probability that this manager is a star now? What's the probability that the manager will beat the market next year, given that he has beaten it in the past three years?

Solution

We first summarize the data by introducing some notations as follows: The chances that a manager will beat the market on the condition that he is a star is:

$$P(B|S) = 0.75 = \frac{3}{4}$$

The chances of a non-star manager beating the market are:

$$P(B|\bar{S}) = 0.5 = \frac{1}{2}$$

The chances of the new manager being a star during the particular time he was added to the analyst's portfolio are exactly the chances that any manager will be made a star, which is unconditional:

$$P[S] = 0.16 = \frac{4}{25}$$

To evaluate the likelihood of him being a star at present, we compute the likelihood of him being a star given that he has beaten the market for three consecutive years, $P(S|3B)$, using the Bayes' theorem:

$$P(S|3B) = \frac{P(3B|S) \times P[S]}{P[3B]}$$

$$P(3B|S) = (\frac{3}{4})^3 = \frac{27}{64}$$

The unconditional chances that the manager will beat the market for three years is the denominator.

$$P[3B] = P(3B|S) \times P[S] + P(3B|\bar{S}) \times P[\bar{S}]$$

$$P[3B] = (\frac{3}{4})^3 \times \frac{4}{25} + (\frac{1}{2})^3 \frac{21}{25} = \frac{69}{400}$$

Therefore:

$$P(S|3B) = \frac{(\frac{27}{64})(\frac{4}{25})}{(\frac{69}{400})} = \frac{9}{23} = 39\%$$

Therefore, there is a 39% chance that the manager will be a star after beating the market for three consecutive years, which happens to be our new belief and is a significant improvement from our old belief, which was 16%.

Finally, we compute the manager's chances of beating the market the next year. This happens to be the summation of the chances of a star beating the market and the chances of a non-star beating the market, weighted by the new belief:

$$P[B] = P(B|S) \times P[S] + P(B|\bar{S}) \times P[\bar{S}]$$

$$P[B] = \frac{3}{4} \times \frac{9}{23} + \frac{1}{2} \times \frac{14}{23} = 60\% = \frac{3}{5}$$

We also have that:

$$P(S|3B) = \frac{P(3B|S) \times P[S]}{P[3B]}$$

The L.H.S of the formula is posterior. The first item on the numerator is the likelihood, and the second part is prior.

Practice Question 1

The probability that the Eurozone economy will grow this year is 18%, and the probability that the European Central Bank (ECB) will loosen its monetary policy is 52%.

Assume that the joint probability that the Eurozone economy will grow and the ECB will loosen its monetary policy is 45%. What is the probability that either the Eurozone economy will grow ***or*** the ECB will loosen its monetary policy?

A. 42.12%

B. 25%

C. 11%

D. 17%

The correct answer is **B**.

The addition rule of probability states that the probability of either of two events happening (E or M in this case) is equal to the sum of their individual probabilities minus the probability of them both occurring.

Let's break it down: The probability of the Eurozone economy growing (E) is given as 18%, or 0.18 in decimal form. The probability of the European Central Bank (ECB) loosening its monetary policy (M) is given as 52%, or 0.52 in decimal form. The joint probability of both these events occurring (EM) is given as 45%, or 0.45 in decimal form.

Using the addition rule of probability, we can calculate the probability of either the Eurozone economy growing or the ECB loosening its monetary policy (E or M) as follows: p(E or M) = p(E) + p(M) - p(EM) = 0.18 + 0.52 - 0.45 = 0.25 or 25%.

Things to Remember

- The addition rule of probability is used when determining the probability of either one event or another occurring.

- This rule states that the probability of either event A or B happening is equal to the sum of their individual probabilities minus the joint probability (if they are not mutually exclusive).

- In this context, it's important to understand that 'or' in probability theory means and/or. It includes scenarios where both events occur simultaneously.

- Joint probabilities are crucial in understanding how two events can interact with each other. They represent the likelihood of two independent events happening at the same time.

- Misinterpretation often occurs when candidates confuse independent and mutually exclusive events. If two events are mutually exclusive, they cannot happen at the same time, which is not our case here.

Practice Question 2

A mathematician has given you the following conditional probabilities:

p(O|T) = 0.62	Conditional probability of reaching the office if the train arrives on time
p(O|T c) = 0.47	Conditional probability of reaching the office if the train does not arrive on time
p(T) = 0.65	Unconditional probability of the train arriving on time
p(O) = ?	Unconditional probability of reaching the office

What is the unconditional probability of reaching the office, p(O)?

A. 0.4325

B. 0.5675

C. 0.3856

D. 0.5244

The correct answer is **B**.

The total probability rule is a fundamental rule in statistics that provides a way to calculate the probability of an event from the probabilities of events in a partition of the sample space. In this case, the events are "the train arrives on time" and "the train does not arrive on time".

The unconditional probability of the train arriving on time, p(T), is given as 0.65. This means that out of all possible outcomes, the train arrives on time 65% of the time. The unconditional probability of the train not arriving on time, p(T c), is the complement of p(T), which is 1 - p(T) = 1 - 0.65 = 0.35. This means that out of all possible outcomes, the train does not arrive on time 35% of the time.

The conditional probability of reaching the office if the train arrives on time, p(O|T), is given as 0.62. This means that if the train arrives on time, the probability of reaching the office is 62%. The conditional probability of reaching the office if the train does not arrive on time, p(O|T c), is given as 0.47. This means that if the train does not arrive on time, the probability of reaching the office is 47%.

Using the total probability rule, we can calculate the unconditional probability of reaching the office, p(O), as follows:

$$\begin{aligned} p(O) &= p(O|T) * p(T) + p(O|Tc) * p(Tc) \\ &= 0.62 * 0.65 + 0.47 * 0.35 \\ &= 0.5675 \end{aligned}$$

Things to Remember

- The total probability rule is a fundamental concept in probability theory, used to calculate the unconditional probability of an event by considering all possible outcomes.
- Conditional probabilities (p(O|T) and p(O|Tc)) are probabilities of an event

occurring given that another event has already occurred. They play a crucial role in calculating the total probability.

- The unconditional probabilities (p(T) and p(Tc)) represent the likelihood of an event without any conditions. The sum of these probabilities is always 1.
- In real-world applications, understanding these concepts can help in risk assessment, decision making under uncertainty, and various other fields where probabilistic models are used.

Practice Question 3

Suppose you are an equity analyst for the XYZ investment bank. You use historical data to categorize the managers as excellent or average. Excellent managers outperform the market 70% of the time and average managers outperform the market only 40% of the time. Furthermore, 20% of all fund managers are excellent managers and 80% are simply average. The probability of a manager outperforming the market in any given year is independent of their performance in any other year.

A new fund manager started three years ago and outperformed the market all three years. What’s the probability that the manager is excellent?

A. 29.53%

B. 12.56%

C. 57.26%

D. 30.21%

The correct answer is **C**.

This probability problem can be best understood by creating a probability matrix. The matrix would look like this:

Kind of manager	Probability	Probability of beating market
Excellent	0.2	0.7
Average	0.8	0.4

Let's denote E as the event of an excellent manager, and A as the event of an average manager. The probability of E (P(E)) is 0.2 and the probability of A (P(A)) is 0.8. Let's denote O as the event of outperforming the market. The probability of O given E (P(O|E)) is 0.7 and the probability of O given A (P(O|A)) is 0.4.

We are interested in finding the probability of E given O (P(E|O)). This can be calculated using the formula:

$$
\begin{aligned}
P(E|O) &= \frac{P(O|E) \times P(E)}{P(O|E) \times P(E) + P(O|A) \times P(A)} \\
&= \frac{(0.7^3) \times 0.2}{(0.7^3) \times 0.2 + (0.4^3) \times 0.8} \\
&= 57.26\%
\end{aligned}
$$

The power of three is used to indicate three consecutive years of outperforming the market.

Things to Remember

- The Practice Question is based on the concept of conditional probability, which refers to the probability of an event given that another event has occurred.

- In this context, P(E|O) represents the probability that a manager is excellent given they have outperformed the market. This is calculated using Bayes' theorem.

- Bayes' theorem allows us to update our initial beliefs (prior probabilities), in light of new evidence. In this case, it's used to update our belief about a manager's ability based on their performance over three years.

- The independence assumption means that each year’s performance does not influence or predict future performances. Hence, we can raise the probabilities to the power of 3 for three consecutive years.

- A common misunderstanding is confusing P(E|O) with P(O|E). The former asks for the likelihood a manager is excellent given their market outperformance, while latter asks for likelihood of outperforming given they are an excellent manager.

<u>Reading 2: Random Variables</u>

After completing this reading, you should be able to:

- Describe and distinguish a probability mass function from a cumulative distribution function and explain the relationship between these two.
- Understand and apply the concept of a mathematical expectation of a random variable.
- Describe the four common population moments.
- Explain the differences between a probability mass function and a probability density function.
- Characterize the quantile function and quantile-based estimators.
- Explain the effect of a linear transformation of a random variable on the mean, variance, standard deviation, skewness, kurtosis, median, and interquartile range.

Random Variables

A random variable is a variable whose possible values are outcomes of a random phenomenon. It is a function that maps outcomes of a random process to real values. It can also be termed as the realization of a random process.

Precisely, if ω is an element of a sample space Ω and x is the realization, then $X(\omega) = x$. Conventionally, random variables are given in upper case (such as X, Y, and Z) while the realized random values are represented in lower case (such as x, y, and z)

For example, let X be the random variable as a result of rolling a die. Therefore, x is the outcome of one roll, and it could take any of the values 1, 2, 3, 4, 5, or 6. The probability that the resulting random variable is equal to 3 can be expressed as:

$$P(X = x) \text{ where } x = 3$$