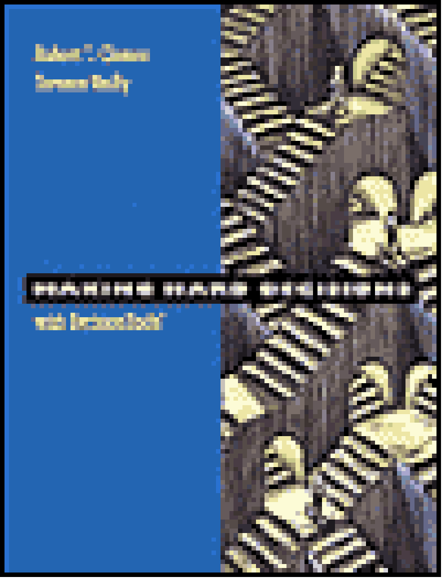

Making Hard Decisions

R. T. Clemen, T. Reilly

Chapter 7

Probability Basics

Draft: Version 1

Introduction

Let A be **an event** with possible outcomes: $A_1, \cdots, A_n$

$$A = \text{“Flipping a coin”}$$

$$A_1 = \{\text{Heads}\} \qquad A_2 = \{\text{Tails}\}$$

The total event Ω (or sample space) of event A is the collection of all possible outcomes of A

$$\Omega = \{\text{Heads, Tails}\}$$

Formally:

$$\Omega = A_1 \cup A_2 \cup \cdots \cup A_{n-1} \cup A_n = \bigcup_{i=1}^{n} A_i$$

Draft: Version 1

Probability rules may be derived using **VENN DIAGRAMS**

1. Probabilities must be **between 0 and 1** for all possible outcomes in the sample space Ω :

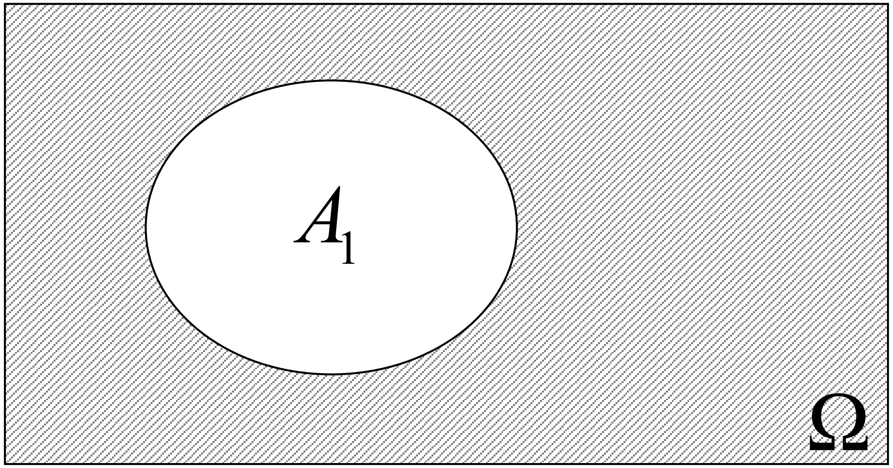

$$0 \leq \Pr(A_i) \leq 1,$$ for all outcomes A_i that are in Ω

Ratio of the **area of the oval** and the **area of the total rectangle** can be interpreted as the probability of the event

2. Probabilities must **add up** if both events **cannot occur at the same time**:

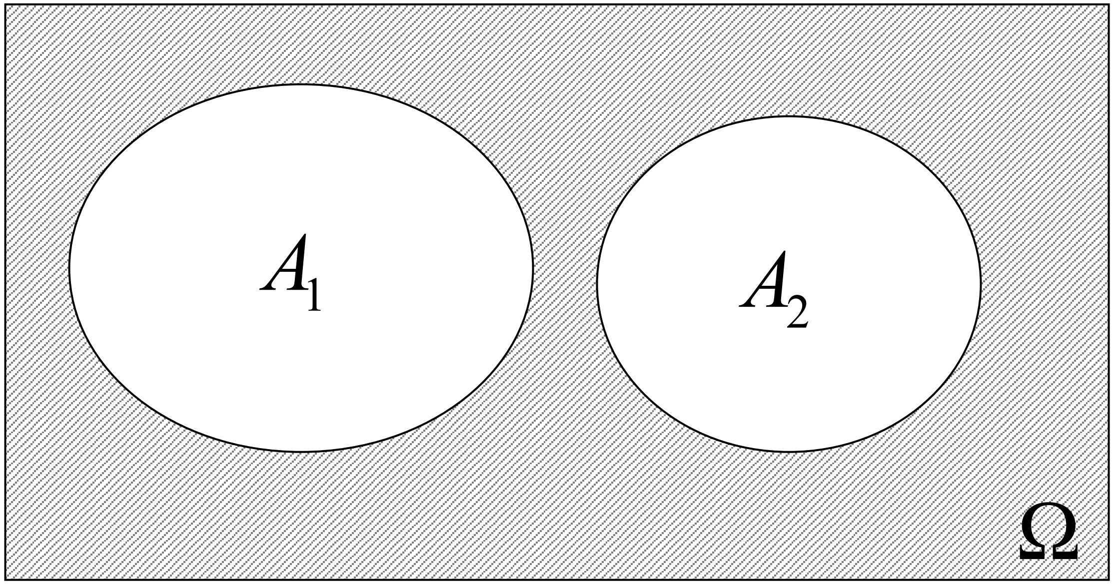

$$A_1 \cap A_2 = \phi \Rightarrow \Pr(A_1 \cup A_2) = \Pr(A_1) + \Pr(A_2)$$

Draft: Version 1

Probability Calculus

3. If $A_1, \cdots, A_n$ are **all the possible outcomes** and **not two of these can occur at the same time**, their **Total Probability** must sum up to 1:

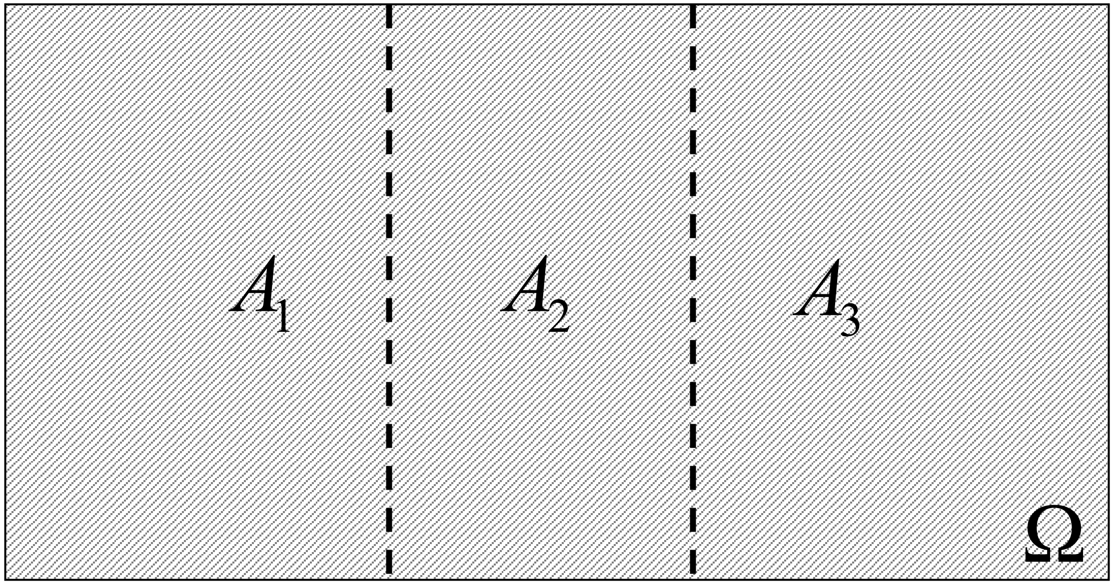

$A_1, \cdots, A_n$ are said to be **collectively exhaustive** and **mutually exclusive**

Draft: Version 1

Probability Calculus

4. The probability of **the complement** of $\overline{A_1}$ equals 1 **minus** the probability of A_1

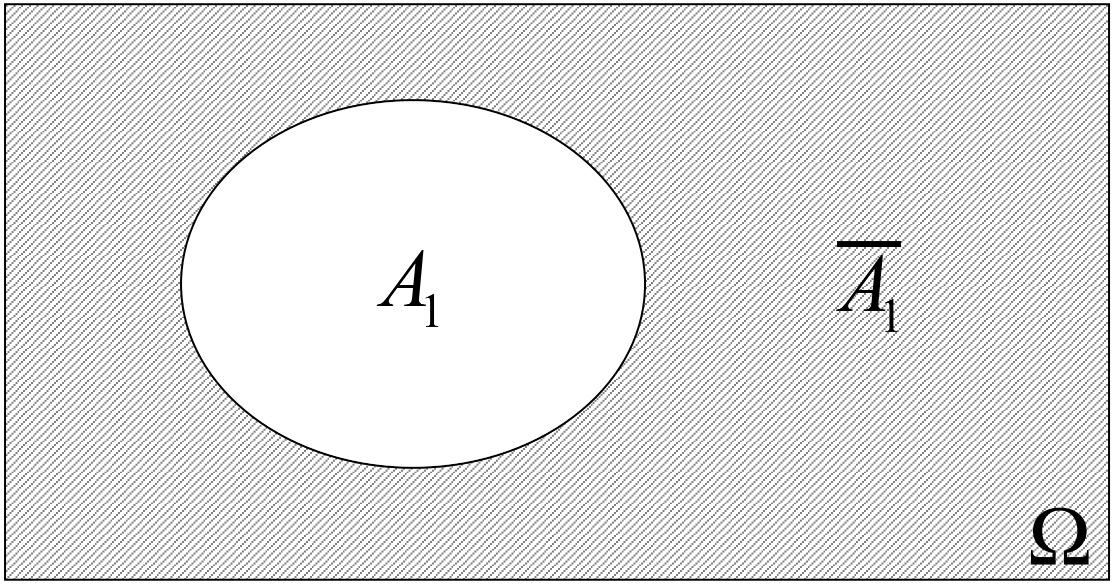

$$\Pr(\overline{A_1}) = 1 - \Pr(A_1)$$

Draft: Version 1

Probability Calculus

5. If two events **can occur at the same time** the probability of **either of them happening or both** equals the **sum** of their individual probability **minus** the probability of them both happening at the same time.

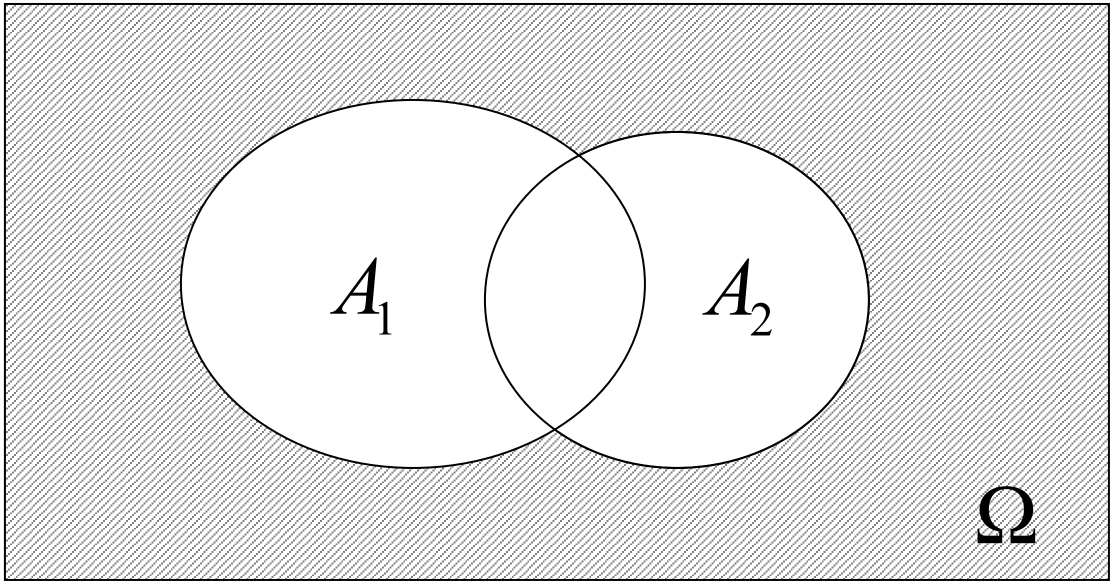

$$\Pr(A_1 \cup A_2) = \Pr(A_1) + \Pr(A_2) - \Pr(A_1 \cap A_2)$$

6. **Conditional probability:**

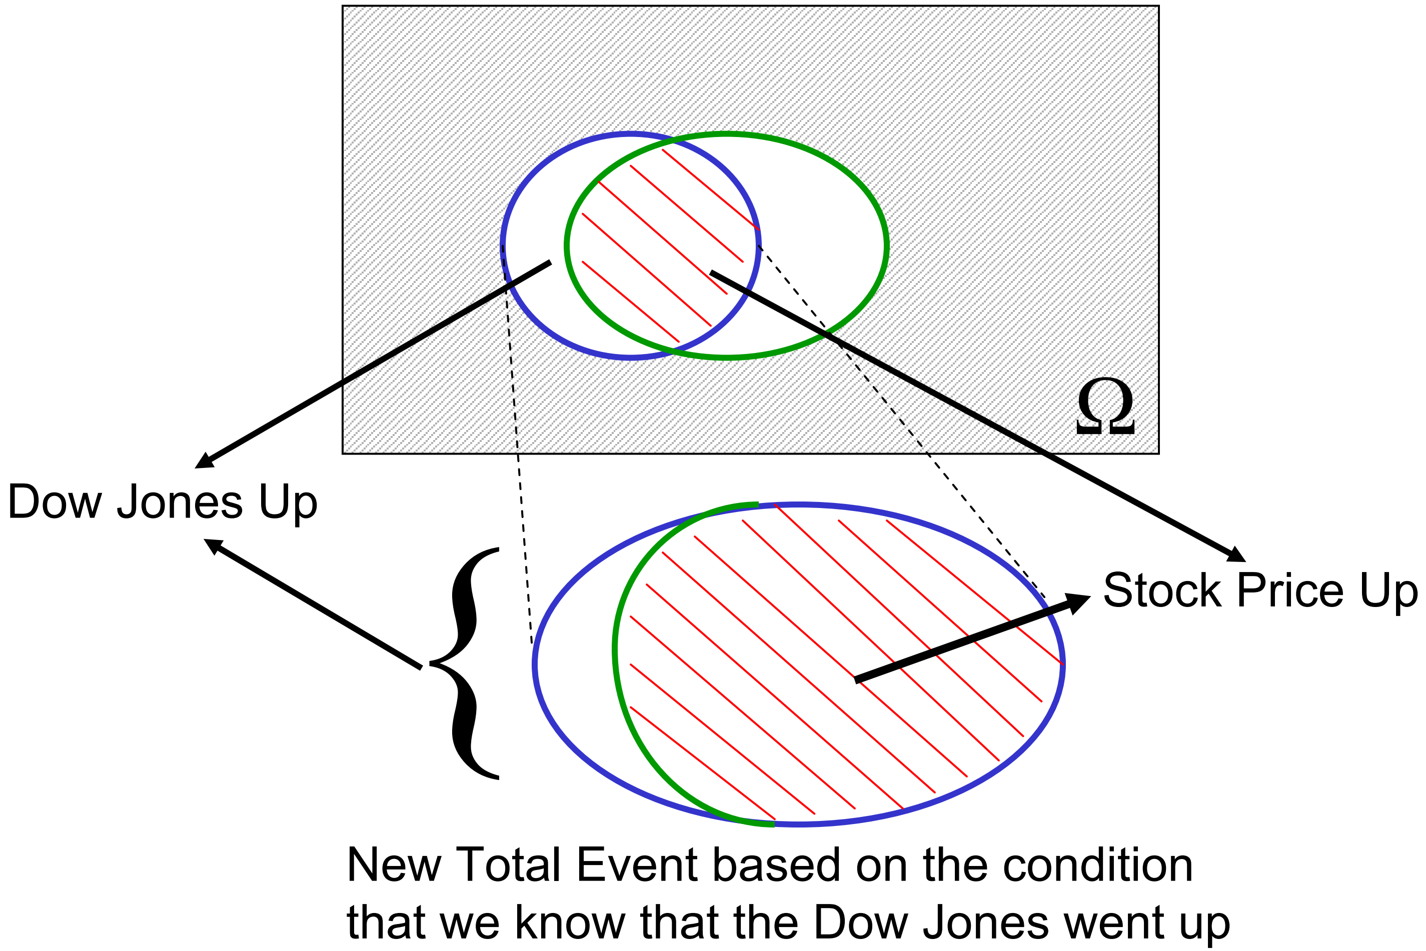

Probability Calculus: Conditional Probability

$$\Pr(Stock\uparrow \mid Dow\uparrow) = \frac{\Pr(Stock\uparrow \cap Dow\uparrow)}{\Pr(Dow\uparrow)}$$

Intuition: If I know that the market as a whole will go up, the chances of the stock of an individual company going up will increase.

$$\Pr(A \mid B) = \frac{\Pr(A \cap B)}{\Pr(B)}$$

Informally: Conditioning on an event coincides with reducing the total event to **the conditioning event**

Probability Calculus: Conditional Probability

Example: The probability of drawing an ace of spades in a deck of 52 cards equals 1/52. However, if I tell you that I have an ace in my hands, the probability of it being the ace of spades equals ¼.

$$\Pr(Spades \mid Ace) = \frac{\Pr(Ace \cap Space)}{\Pr(Ace)} = \frac{1/52}{4/52} = \frac{1}{4}$$

Note also that:

$$\Pr(B \mid A) = \frac{\Pr(A \cap B)}{\Pr(A)}$$

Draft: Version 1

7. **Multiplicative Rule:** Calculating the probability of two events happening at the same time.

$$\Pr(A_i \cap B) = \Pr(B \mid A) * \Pr(A)$$
$$= \Pr(A \mid B) * \Pr(B)$$

8. **Independence between two events:** Informally, two events are independent if information about one does not provide you any information about the other and vice versa. Consider:

Event A with possible outcomes $A_1, \cdots, A_n$

Event B with possible outcomes $B_1, \cdots, B_m$

Draft: Version 1

Probability Calculus: Independence

Example: A is the event of **flipping a coin** and B is the event of **throwing a dice**. If you know the outcome of flipping the coin you do not learn anything about the outcome of throwing the dice (regardless of the outcome of flipping the coin). Hence, these two events are independent.

Formal definition of independence between event A and event B :

$$\Pr(A_i \mid B_j) = \Pr(A_i)$$

For all possible combinations A_i **and** B_j

Probability Calculus: Independence

Equivalent definitions of independence between A event and event B :

1. $$\Pr(B_j \mid A_i) = \Pr(B_j)$$

For all possible combinations A_i and B_j

2. $$\Pr(A_i \cap B_j) = \Pr(A_i) \times \Pr(B_j)$$

For all possible combinations A_i and B_j

Independence/dependence in influence diagrams:

- **No arrow** between two chance nodes implies **independence** between the uncertain events
- **An arrow** from a chance event A to a chance event B does **not mean** that "A causes B". It indicates that information about A helps in determining the likelihood of outcomes of B.

Draft: Version 1

Probability Calculus: Conditional Independence

Example: The performance of a person on any **IQ test** is uncertain and may range anywhere from **0% to 100%.** However, if you to know that the person in question is highly intelligent it is expected his\her score will be **high**, e.g. ranging anywhere from **90% to 100%.**

On the other hand, the person's IQ **does not explain** this remaining uncertainty, and it may be considered **measurement error** affected by other conditions. For example, having a good night sleep during the previous night. On any two IQ tests, these measurement errors may be reasonably modeled as **independent**, **if** we know the IQ of the person.

Probability Calculus: Conditional Independence

Event A with possible outcomes $A_1, \cdots, A_n$

Event B with possible outcomes $B_1, \cdots, B_m$

Event C with possible outcomes $C_1, \cdots, C_p$

Formal definition: Event A and event B are conditionally independent given event C if and only if

$$\Pr(A_i \mid B_j, C_k) = \Pr(A_i \mid C_k)$$

For all possible combinations A_i, B_j and C_k

Informally: If I already know C, any information or knowledge about B does not tell me anything more about A

Draft: Version 1

Probability Calculus: Conditional Independence

Equivalent definitions: Event A and event B are conditionally independent given event C if and only if

1. $$\Pr(B_j \mid A_i, C_k) = \Pr(B_j \mid C_k)$$

For all possible combinations A_i, B_j and C_k

2. $$\Pr(A_i \cap B_j \mid C_k) = \Pr(A_i \mid C_k) \times \Pr(B_j \mid C_k)$$

For all possible combinations A_i, B_j and C_k

Draft: Version 1

Probability Calculus: Conditional Independence

Equivalent definitions: Event A and event B are conditionally independent given event C if and only if

1. $$\Pr(B_j \mid A_i, C_k) = \Pr(B_j \mid C_k)$$

For all possible combinations A_i, B_j and C_k

2. $$\Pr(A_i \cap B_j \mid C_k) = \Pr(A_i \mid C_k) \times \Pr(B_j \mid C_k)$$

For all possible combinations A_i, B_j and C_k

Draft: Version 1

Conditional independence in influence diagrams:

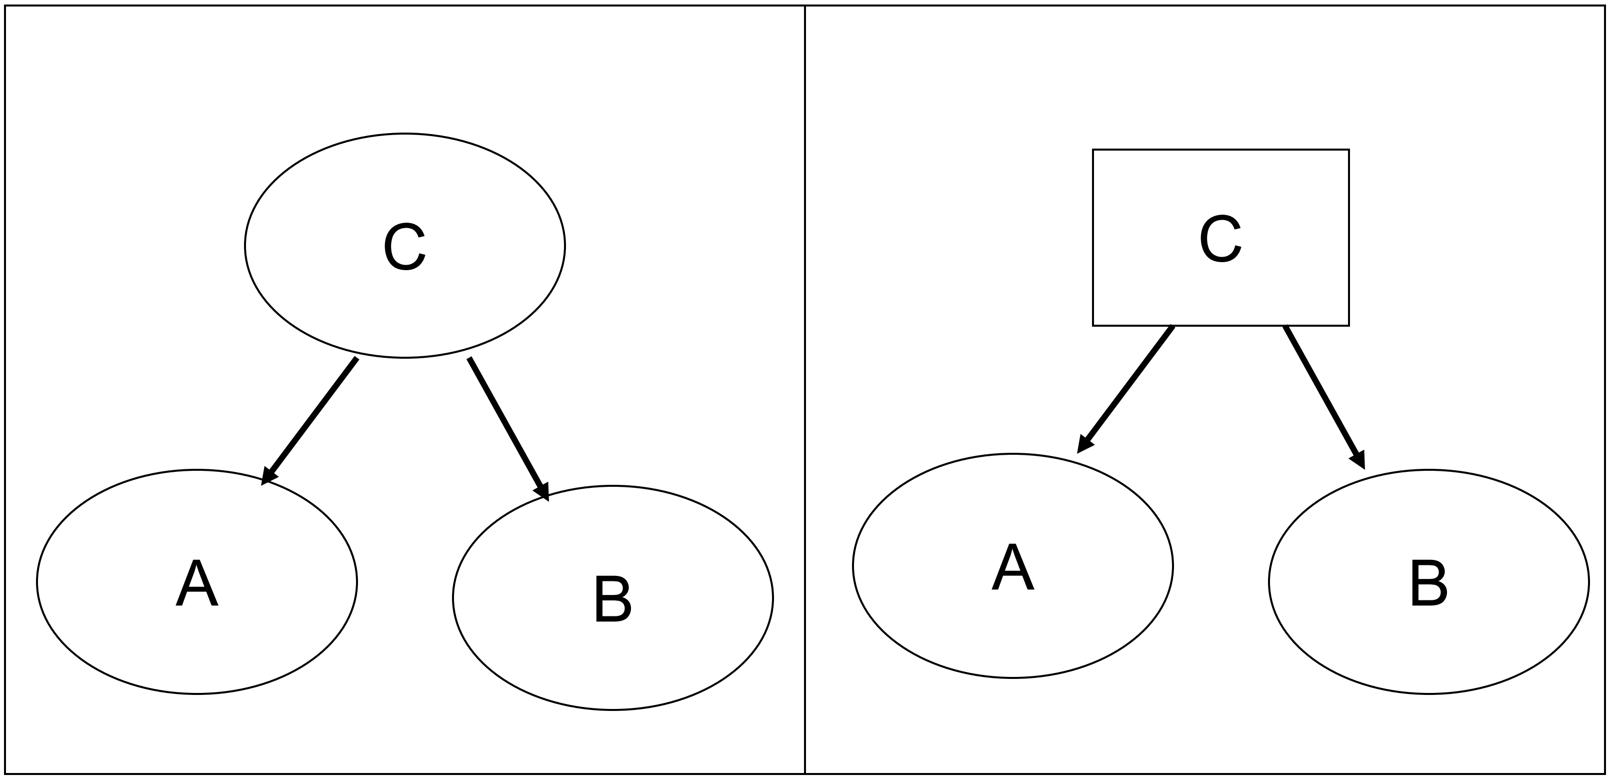

Probability Calculus: Law of Total Probability

- Let $B_1, \cdots, B_3$ be mutually exclusive, collectively exhaustive:

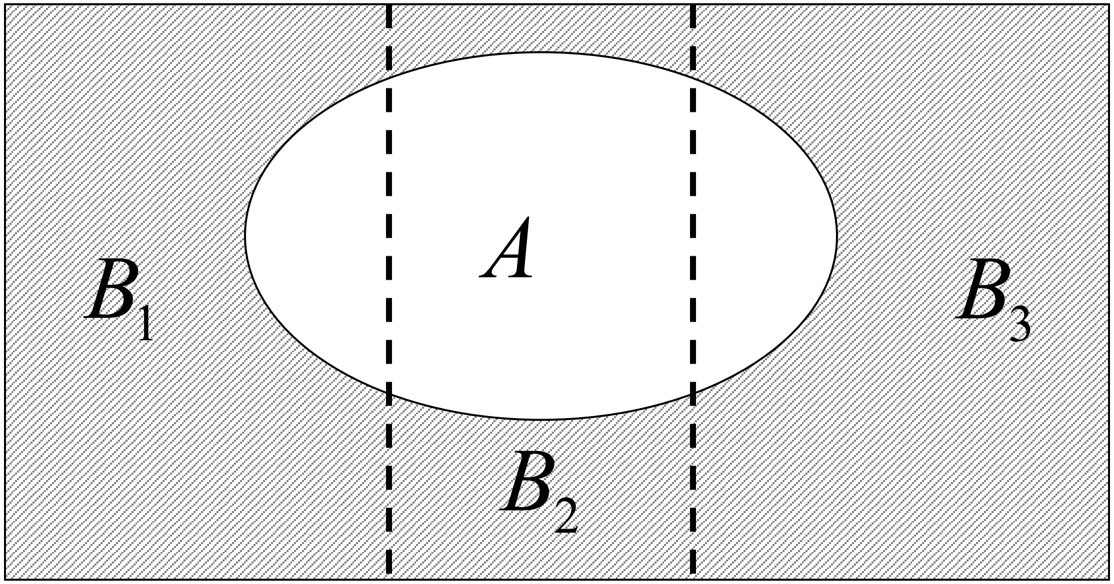

$$\Pr(A) = \Pr(A \cap B_1) + \Pr(A \cap B_2) + \Pr(A \cap B_3) \Leftrightarrow$$

$$\Pr(A) = \Pr(A \mid B_1)\Pr(B_1) + \Pr(A \mid B_2)\Pr(B_2) + \Pr(A \mid B_3)\Pr(B_3)$$

Example:

X = System fails

A = Component A fails,

B = Component B fails,

C = Component C fails

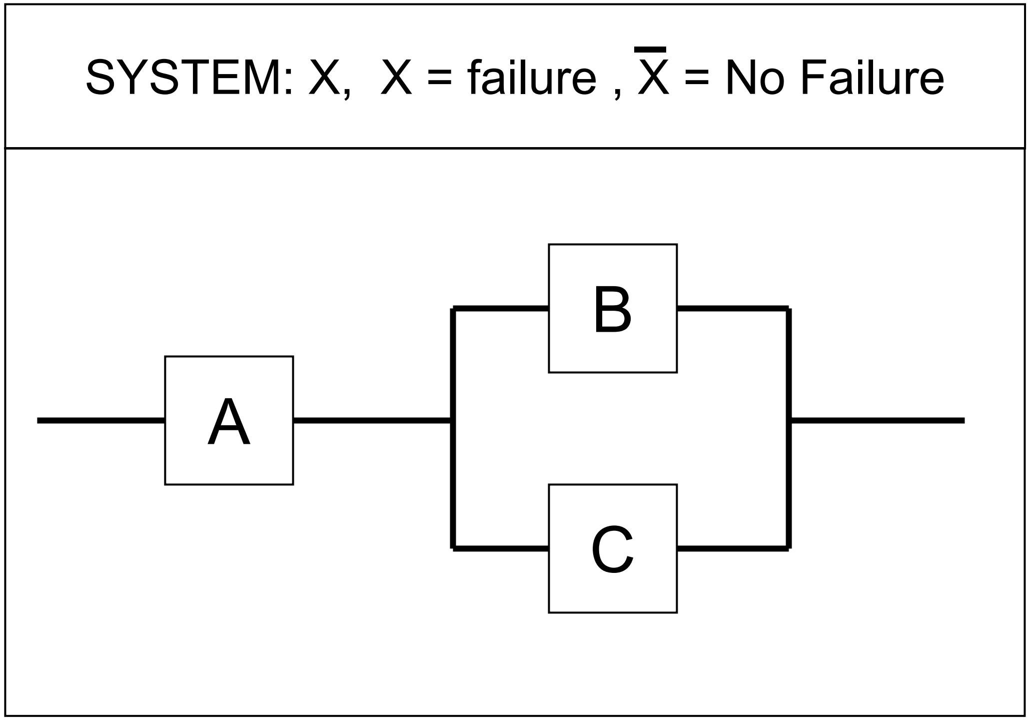

Assume that components A, B and C operate independently.

Probability Calculus: Law of Total Probability

Task:

Write the probability of failure Pr(X) as a function of the component failure probabilities Pr(A), Pr(B) and Pr(C).

$$1.\ \Pr(X) = \Pr(X \mid A)\Pr(A) + \Pr(X \mid \overline{A})\Pr(\overline{A}) =$$
$$= 1 * \Pr(A) + \Pr(X \mid \overline{A})\Pr(\overline{A})$$

$$2.\ \Pr(X \mid \overline{A}) = \Pr(X \mid B, \overline{A})\Pr(B \mid \overline{A}) +$$
$$\Pr(X \mid \overline{B}, \overline{A})\Pr(\overline{B} \mid \overline{A})$$
$$= \Pr(X \mid B, \overline{A})\Pr(B) + 0 * \Pr(\overline{B})$$
$$= \Pr(X \mid B, \overline{A})\Pr(B)$$

Substitute result 2 into 3

$$3.\ \Pr(X) = \Pr(A) + \Pr(X \mid B, \overline{A})\Pr(B)\Pr(\overline{A})$$

Draft: Version 1

Intermediate conclusion: Hence we need to further develop

$$\Pr(X \mid B, \overline{A})$$

$$4. \Pr(X \mid B, \overline{A}) = \Pr(X \mid C, B, \overline{A}) \Pr(C \mid B, \overline{A}) + \Pr(X \mid \overline{C}, B, \overline{A}) \Pr(\overline{C} \mid B, \overline{A})$$

$$= 1 * \Pr(C) + 0 * \Pr(\overline{C}) = \Pr(C)$$

Substitute result 4 into 3

$$5. \Pr(X) = \Pr(A) + \Pr(C)\Pr(B)\Pr(\overline{A})$$

$$6. \Pr(\overline{A}) = 1 - \Pr(A)$$ **Substitute result 6 into 5**

$$7. \Pr(X) = \Pr(A) + \Pr(C)\Pr(B) - \Pr(C)\Pr(B)\Pr(A)$$

Draft: Version 1

Example: Oil Wildcatter Problem

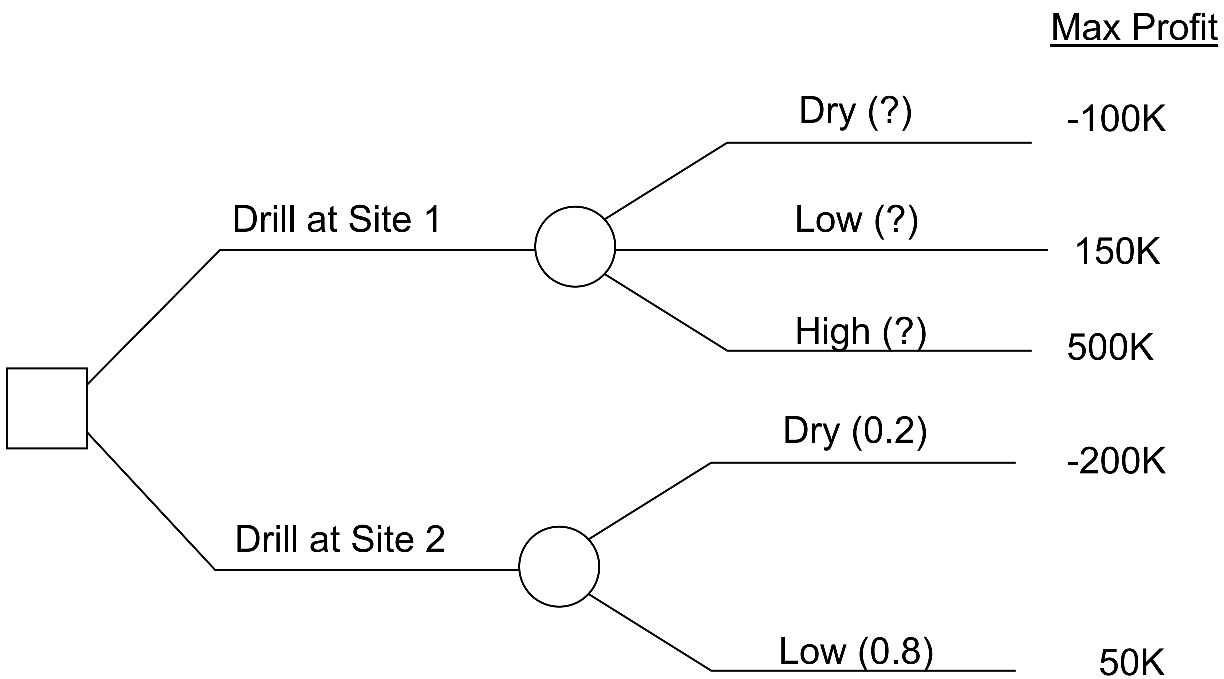

Payoff at site 1 is uncertain. **Dominating factor** in eventual payoff at Site 1 is the presence of a dome or not.

Probability Calculus: Law of Total Probability

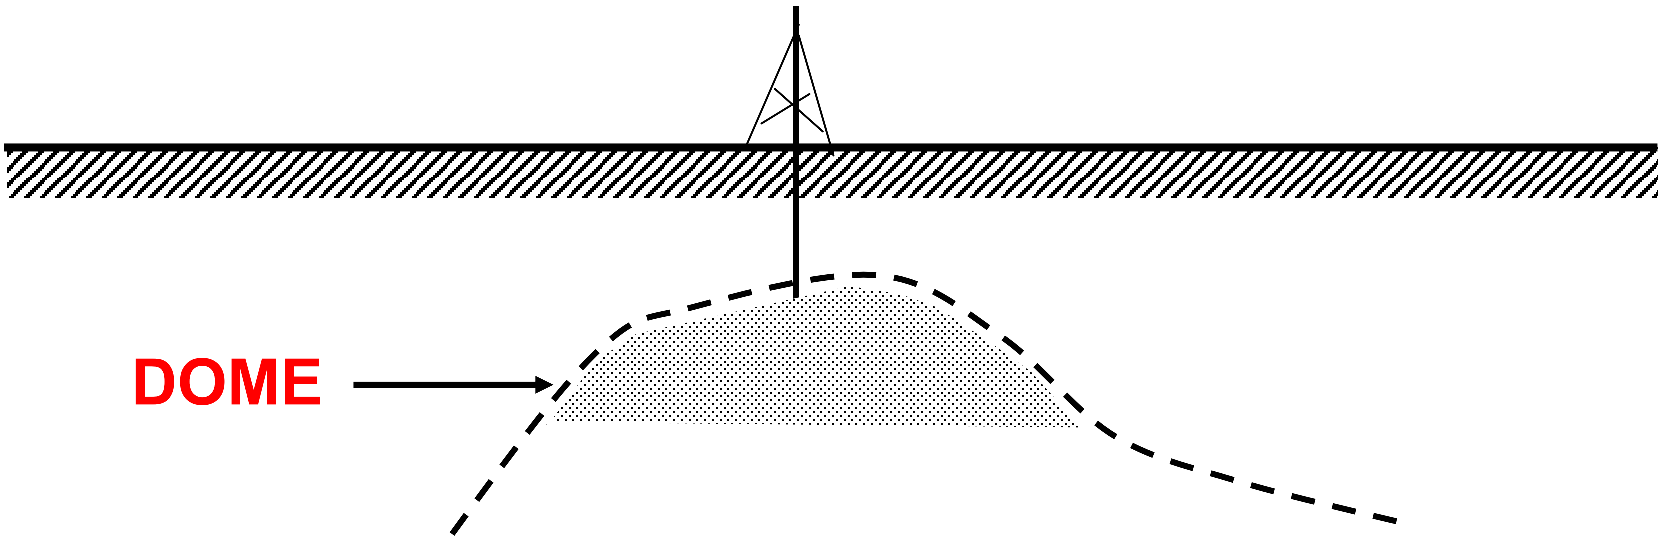

Pr(Dome)	Pr(No Dome)
0.600	0.400

Outcome	Pr(Outcome|Dome)
Dry	0.600
Low	0.250
High	0.150

Outcome	Pr(Outcome|No Dome)
Dry	0.850
Low	0.125
High	0.025

Draft: Version 1

Probability Calculus: Law of Total Probability

$$\Pr(Dry) = \Pr(Dry \mid Dome)\Pr(Dome) + \Pr(Dry \mid No\ Dome)\Pr(No\ Dome)$$
$$= 0.600 * 0.600 + 0.850 * 0.400 = 0.700$$

$$\Pr(Low) = \Pr(Low \mid Dome)\Pr(Dome) + \Pr(Low \mid No\ Dome)\Pr(No\ Dome)$$
$$= 0.250 * 0.600 + 0.125 * 0.400 = 0.200$$

$$\Pr(High) = \Pr(High \mid Dome)\Pr(Dome) + \Pr(High \mid No\ Dome)\Pr(No\ Dome)$$
$$= 0.150 * 0.600 + 0.025 * 0.400 = 0.100$$

Probability Calculus: Law of Total Probability

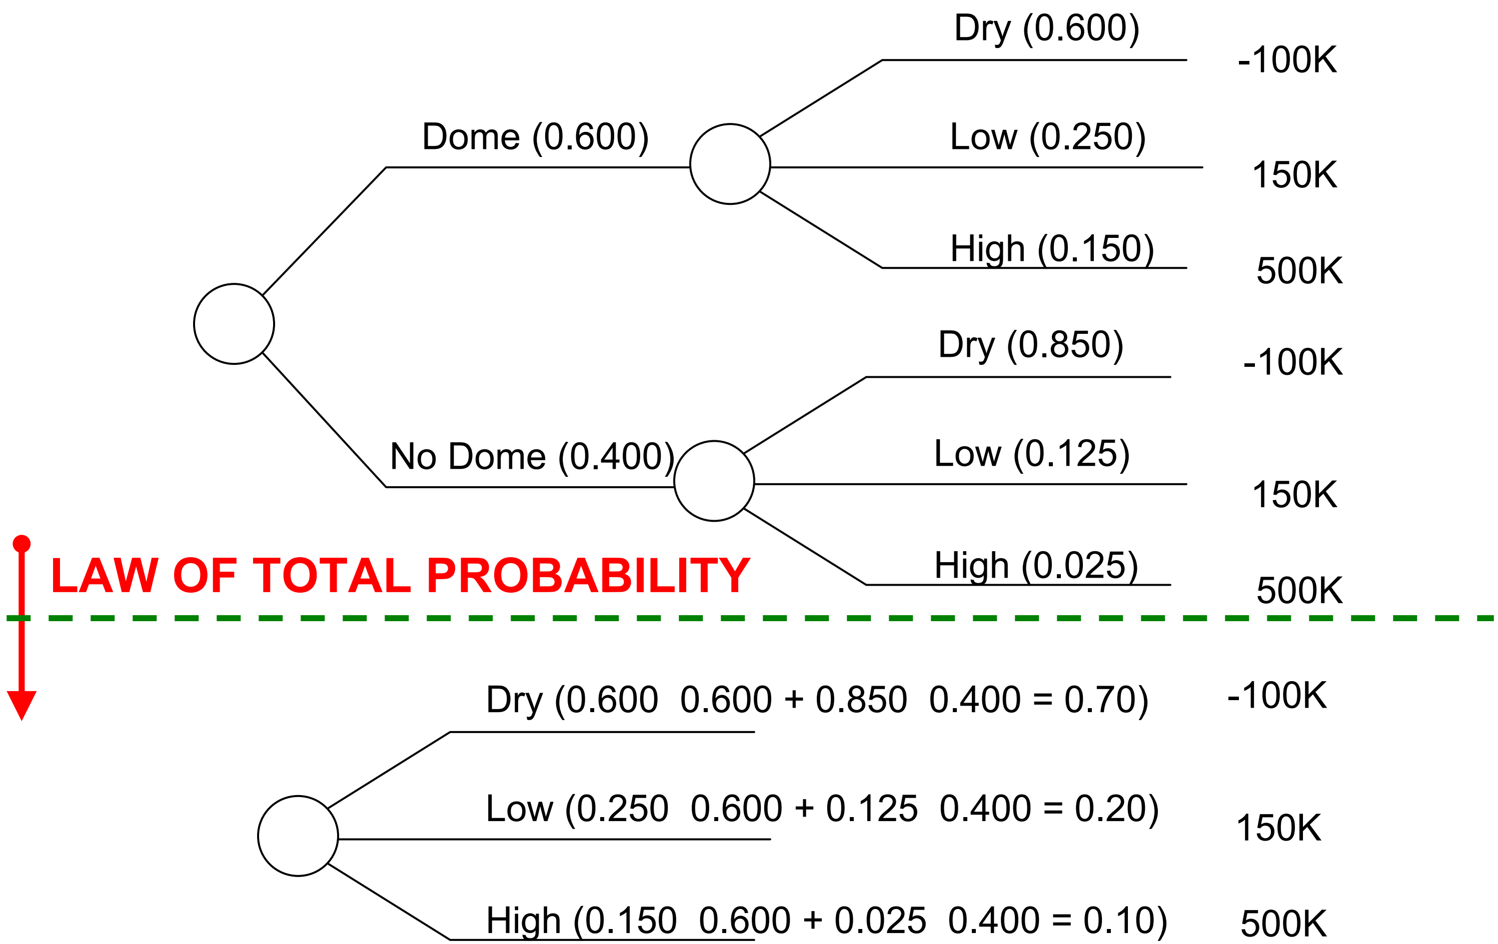

Informally: when we apply **LOTP** we are collapsing a probability tree

Draft: Version 1

Probability Calculus: Bayes Theorem

- Let $B_1, \cdots, B_3$ be mutually exclusive, collectively exhaustive:

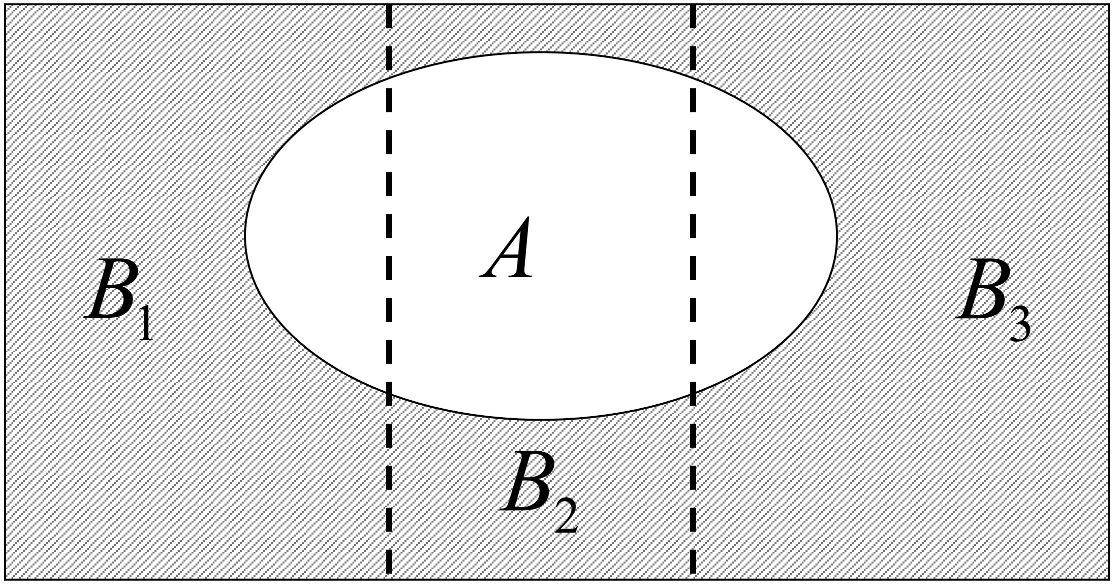

1. From **the multiplicative rule** it follows that:

$$\Pr(A \cap B_j) = \Pr(B_j \mid A)\Pr(A) = \Pr(A \mid B_j)\Pr(B_j)$$

Draft: Version 1

Probability Calculus: Bayes Theorem

2. Dividing the **LHS** and **RHS** of 1. by Pr(A) yields:

$$\Pr(B_j \mid A) = \frac{\Pr(A \mid B_j)\Pr(B_j)}{\Pr(A)}$$

3. We may rewrite Pr(A) using **the Law of Total Probability**, yielding

$$\Pr(A) = \Pr(A \mid B_1)\Pr(B_1) + \Pr(A \mid B_2)\Pr(B_2) + \Pr(A \mid B_3)\Pr(B_3)$$

4. Substituting the result of 3. into 2. gives perhaps the most well known theorem in probability theory:

Bayes Theorem.

Probability Calculus: Bayes Theorem

$$\Pr(B_j \mid A) = \frac{\Pr(A \mid B_j)\Pr(B_j)}{\Pr(A \mid B_1)\Pr(B_1) + \Pr(A \mid B_2)\Pr(B_2) + \Pr(A \mid B_3)\Pr(B_3)}$$

Thomas Bayes (1702 to 1761):

Bayes' theory of probability was published in 1764. His conclusions were accepted by Laplace in 1781, rediscovered by Condorcet, and remained unchallenged until Boole questioned them. Since then Bayes' techniques have been subject to controversy.

Source: http://www-gap.dcs.st-and.ac.uk/~history/Mathematicians/Bayes.html

Draft: Version 1

Probability Calculus: Bayes Theorem

Oil Wildcatter Problem Example Continued:

- We drilled at site 1 and the well is a high producer. Given this new information what are the chances that a dome exists? (Perhaps that information is important when attracting additional investors.)

1. From **the rule for conditional probability** it follows that:

$$\Pr(\textit{Dome} \mid \textit{High}) = \frac{\Pr(\textit{High} \mid \textit{Dome})\Pr(\textit{Dome})}{\Pr(\textit{High})}$$

2. From **the LOTP** it follows that:

$$\Pr(\textit{High}) = \Pr(\textit{High} \mid \textit{Dome})\Pr(\textit{Dome}) + \Pr(\textit{High} \mid \textit{No Dome})\Pr(\textit{No Dome})$$

Draft: Version 1

Probability Calculus: Bayes Theorem

Oil Wildcatter Problem Example Continued:

3. Substitution of 2 in 1 yields:

$$\Pr(\boldsymbol{Dome} \mid \boldsymbol{High}) =$$

$$\frac{\Pr(\boldsymbol{High}|\boldsymbol{Dome})\Pr(\boldsymbol{Dome})}{\Pr(\boldsymbol{High}|\boldsymbol{Dome})\Pr(\boldsymbol{Dome})+\Pr(\boldsymbol{High}|\boldsymbol{No\ Dome})\Pr(\boldsymbol{No\ Dome})} =$$

$$\frac{0.150*0.600}{0.150*0.600+0.0250*0.400} = 0.90$$

Pr(Dome) – **The Prior Probability**

Pr(Dome|Data) – **The Posterior Probability**

Data = "The well is a high produces"

Draft: Version 1

Oil Wildcatter Problem Example Continued:

- Notice that Pr(Dry),Pr(Low) and Pr(High) have been inserted in the tree.These were calculated using **LOTP**.

- Notice that Pr(Dome|High) has been inserted as well This one was calculated using **Bayes Theorem.**

- We need to fill out the Remainder of the question Marks.

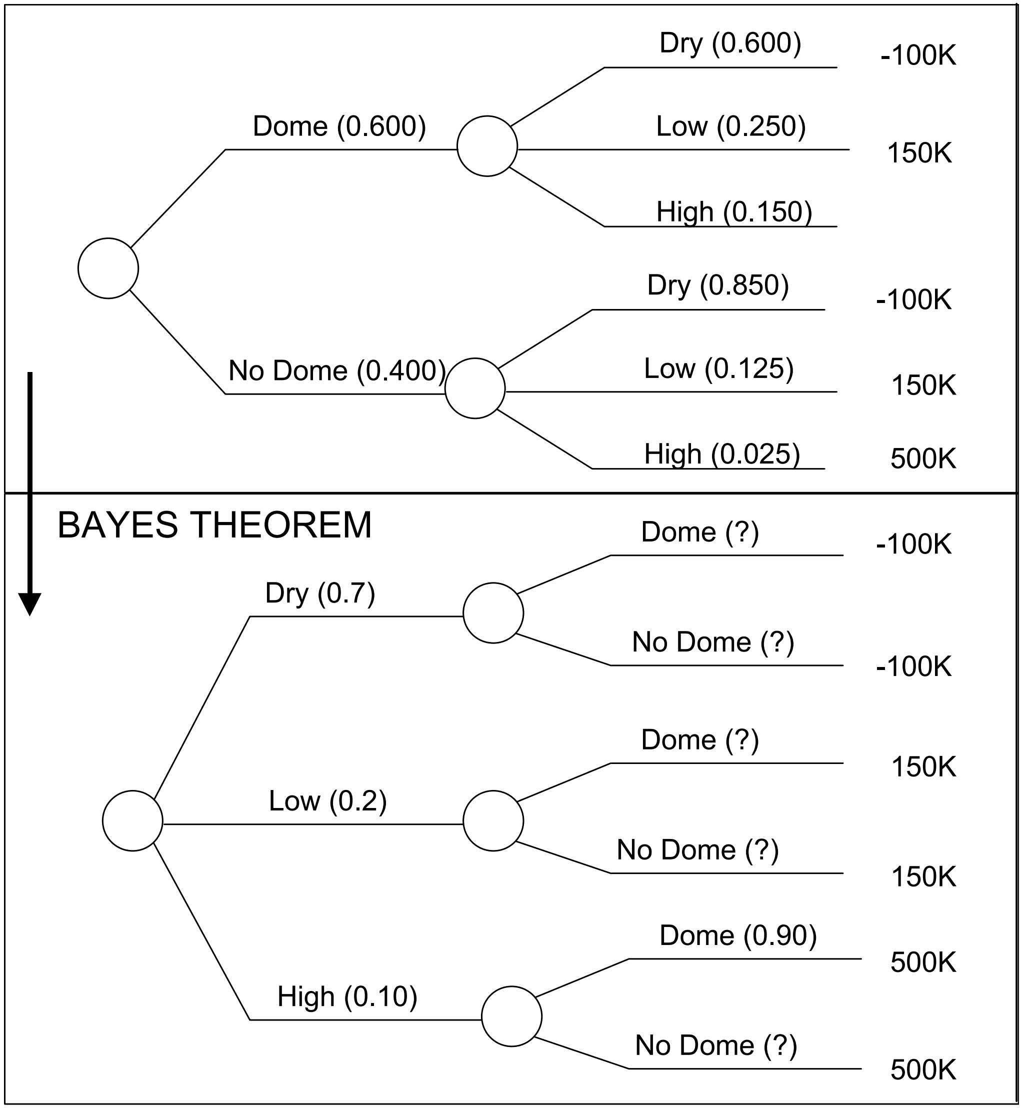

Probability Calculus: Bayes Theorem

Oil Wildcatter Problem Example Continued:

When we **reverse the order of the chance nodes** in a decision tree we need to apply **Bayes Theorem**

	Pr(Dome)	Pr(No Dome)						
	0.600	0.400						
X	**Pr(X|Dome)**	**Pr(X|No Dome)**	**Pr(X $\cap$ Dome)**	**Pr(X $\cap$ No Dome)**	**Pr(X)**	**Pr(Dome|X)**	**Pr(No Dome|X)**	**Check**
Dry	0.600	0.850	0.360	0.340	0.700	0.514	0.486	1.000
Low	0.250	0.125	0.150	0.050	0.200	0.750	0.250	1.000
High	0.150	0.025	0.090	0.010	0.100	0.900	0.100	1.000
Check	1.000	1.000	0.600	0.400	1.000			

Draft: Version 1

Next, we allocate the probabilities from the table at their appropriate locations in the tree

Probability Calculus: Bayes Theorem

Oil Wildcatter Problem Example Continued:

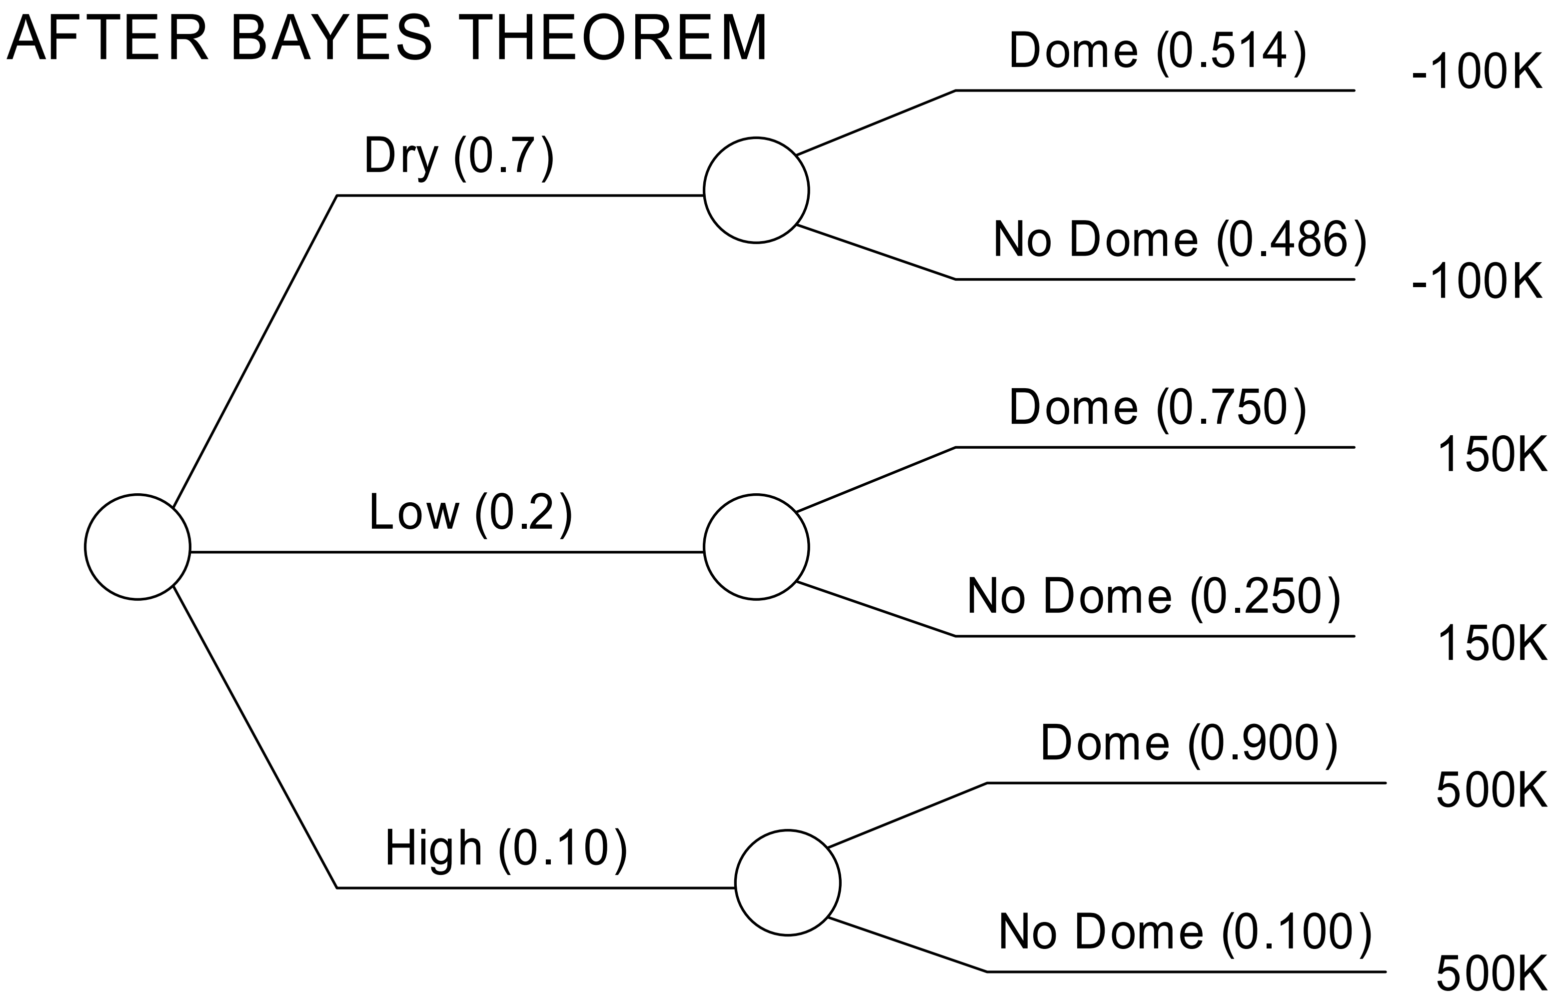

The Game Show Example:

Suppose we have a game show host and you. There are **three doors** and one of them contains a prize. The game show host knows the door containing the prize but of course does not convey this information to you. He asks you to pick a door. **You picked Door 1** and are walking up to door 1 to open it when the game show host screams: **STOP!**.

You stop and the game show host **shows Door 3** which appears to be empty. Next, the game show asks:

"DO YOU WANT TO SWITCH TO DOOR 2?"

WHAT SHOULD YOU DO?

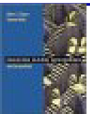

The Game Show Example:

Assumption 1: The game show host will never show the door with the prize.

Assumption 2: The game show will never show the door that you picked.

Define:

D_i ={Prize is behind door i }, i=1,…,3

H_i ={Host shows **Door i** containing no prize **after you selected Door 1**}, i=1,…,3

1. It seems reasonable to set prior probabilities: $\Pr(D_i) = \frac{1}{3}$

Probability Calculus: Bayes Theorem

2. Apply LOTP to calculate $\Pr(H_3)$:

$$\Pr(\boldsymbol{H}_3) = \sum_{\boldsymbol{i}=1}^{3} \Pr(\boldsymbol{H}_3 \mid \boldsymbol{D}_{\boldsymbol{i}}) \Pr(\boldsymbol{D}_{\boldsymbol{i}}) =$$

$$= \frac{1}{2} * \frac{1}{3} + 1 * \frac{1}{3} + 0 * \frac{1}{3} = \frac{1}{2}$$

3. Calculate $\Pr(D_1|H_3)$:

$$\Pr(D_1 \mid H_3) = \frac{\Pr(H_3 \mid D_1)\Pr(D_1)}{\Pr(H_3)} = \frac{\frac{1}{2} * \frac{1}{3}}{\frac{1}{2}} = \frac{1}{3}$$

4. Apply the complement rule:

$$\Pr(D_2 \mid H_3) = 1 - \Pr(D_1 \mid H_3) = 1 - \frac{1}{3} = \frac{2}{3}$$

So YES, you should **SWITCH** since
you would **increase your chances of winning!**

Draft: Version 1

Probability Calculus: Uncertain Quantities

Example: When a student attempts to log on to a computer time-sharing system, either all ports are busy (B) in which case the student will fail to obtain access, or else there is at least one port free (F), in which case the student will be successful in accessing the system.

Total Event: $\Omega = \{B, F\}$

Definition: For a given total event Ω , **a random variable** (rv) is any rule that associates a number with each Ω outcome in.

In mathematical language, a random variable is a function whose domain is the total event and whose range is the real numbers.

Draft: Version 1

Probability Calculus: Uncertain Quantities

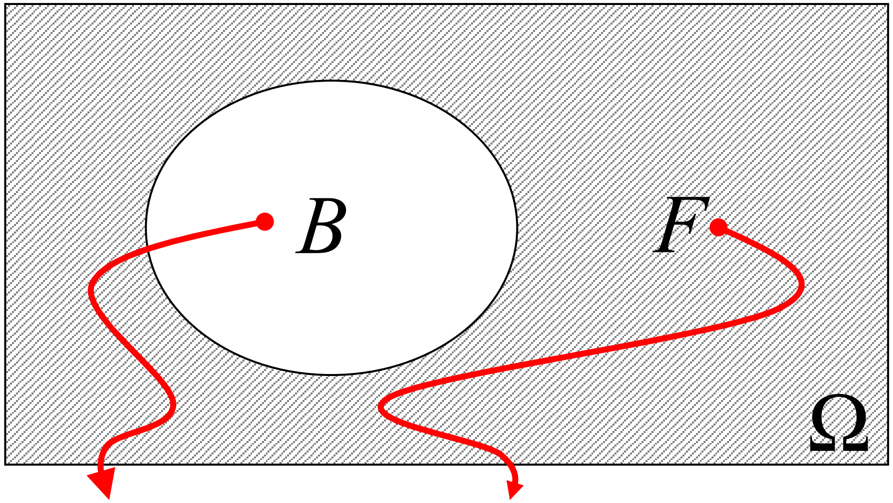

$$X(B) = 0 \qquad X(F) = X(\bar{B}) = 1$$

Example: Consider the experiment in which batteries are examined until a good (G) is obtained. Let B denote a bad Battery.

Total Event: $\Omega = \{G, BG, BBG, BBBG, \ldots\}$

Draft: Version 1

Define a rv X as follows:

X = The **number of batteries** examined before the experiment terminates.

Then:

$$X(G)=1, X(BG)=2, X(BBG)=3, etc$$

The **argument** of the random variable function is typically omitted. Hence, one writes

$$\Pr(X=2)=\Pr(Second\ \ Battery\ \ Works)$$

Note that the above statement **only has meaning** with the above definition of the random variable. It is **good practice** to always include **the definition of a random variable in words.**

The nature of random variables can be **discrete** and **continuous.**

Definition:

A discrete random variable is an rv whose possible values either constitute a finite set or else can be listed in an infinite sequence in which there is a first element, a second element, and so on. **(Think of the previous batter example).**

A random variable is continuous if its set of possible values consists of an entire interval on the number line. **(For example, the failure time of a component).**

Draft: Version 1

Discrete Probability Distributions

The nature of random variables can be
discrete and **continuous.**

Definition:
A discrete random variable is an rv whose possible values either constitute a finite set or else can be listed in an infinite sequence in which there is a first element, a second element, and so on. **(Think of the previous batter example).**

A random variable is continuous if its set of possible values consists of an entire interval on the number line. **(For example, the failure time of a component).**

Discrete Probability Distributions

Example:

Y = Number of Raisins in an Oatmeal Cookie

Assume possible outcomes: Y=1,…,5

Definition: The probability mass function (PMF) of Y is the collection of probabilities such that $\Pr(Y=i)=p_i$

Y	1	2	3	4	5
Pr(Y=i)	0.1	0.15	0.3	0.35	0.1

Note that:

$$\Pr(\boldsymbol{Y}=1)+\Pr(\boldsymbol{Y}=2)+\Pr(\boldsymbol{Y}=3)+\Pr(\boldsymbol{Y}=4)+\Pr(\boldsymbol{Y}=5)=$$

$$\sum_{\boldsymbol{i}=1}^{5}\Pr(\boldsymbol{Y}=\boldsymbol{i})=1$$

Draft: Version 1

Graphical Depictions of PMF's

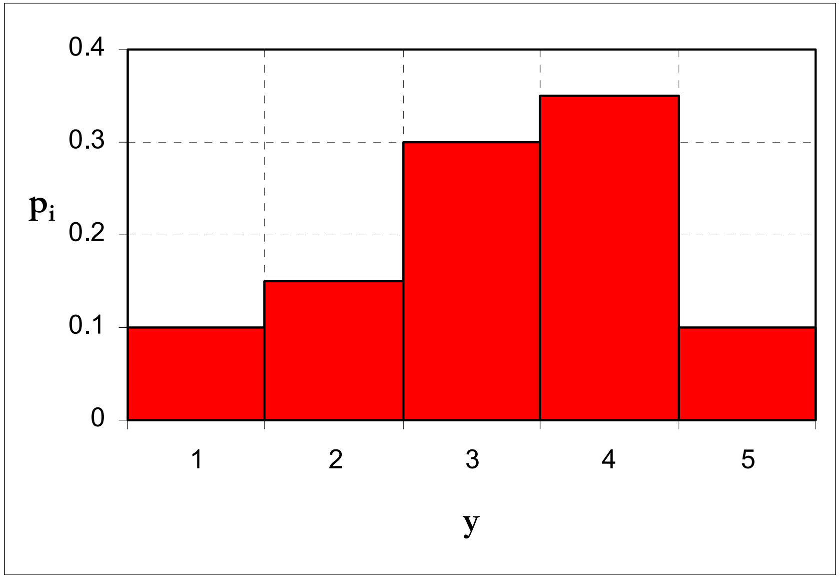

A Histogram

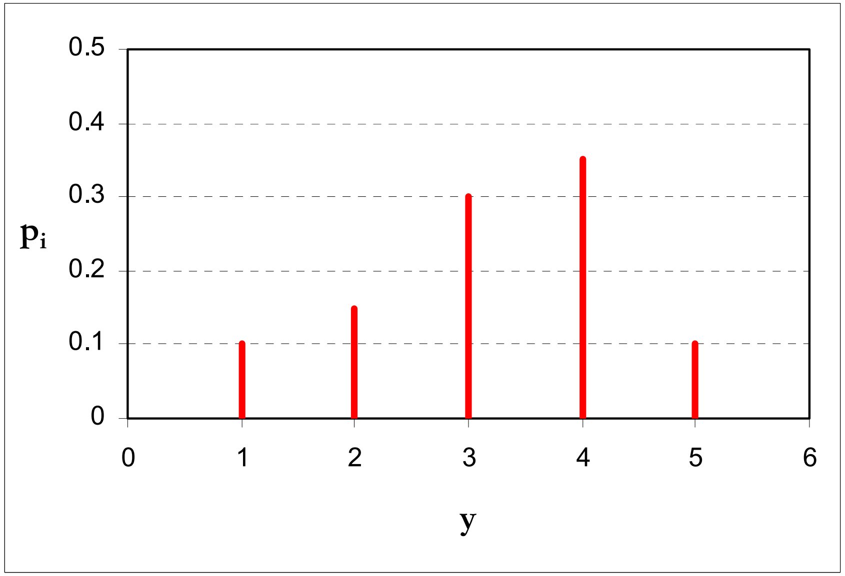

An Line Graph

Definition: The **cumulative distribution function** (CDF) of Y at y is the sum of the probabilities such that Y≤ y:

$$F(y) = \Pr(Y \le y) = \sum_{i:i \le y} \Pr(Y = i)$$

Draft: Version 1

Graphical Depictions of CDF:

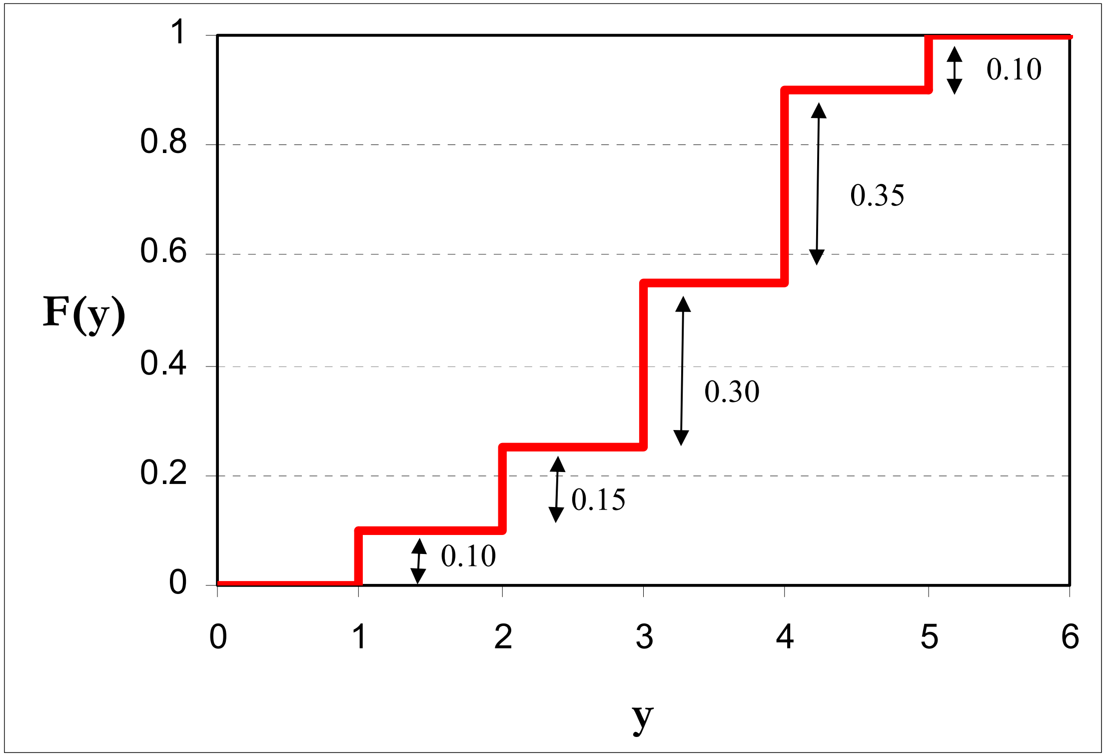

In Decision Analysis a **CDF** is referred to as a **CUMMULATIVE RISK PROFILE** and a **PMF** is referred to as a **RISK PROFILE**

Probability Calculus: Expected Value

We know the random variable Y has many possible outcomes. However, if you were forced to give a **"BEST GUESS"** for Y, what number would you give? (Managers, CEO's, Senators etc. typically like POINT ESTIMATES, unfortunately). Why not use the expected value of Y?

$$E\,[Y] = \sum_{i=1}^{n} y_i \times \Pr(Y = y_i) = \sum_{i=1}^{n} y_i \times p_i$$

Interpretation: If you were able to observe many outcomes of Y, the calculated average of all the outcomes would be close to E[Y].

Draft: Version 1

- If Z = g(Y):

$$E[Z] = \sum_{i=1}^{n} g(y_i) \times \Pr(Y = y_i) = \sum_{i=1}^{n} g(y_i) \times p_i$$

Probability Calculus: Expected Value

- If Z=aY+b, a,b constants, Y a rv: $E[Z]=aE[Y]+b$
- If Z=aX+bY, a,b const., X,Y a rv: $E[Z]=aE[X]+bE[Y]$

Oatmeal Cookie Example:

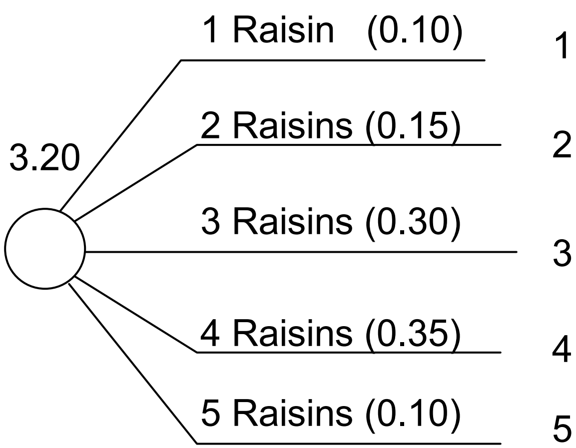

#Raisins	# Raisins*Pr(Y = # Raisins)
1	1*0.10=0.10
2	2*0.15=0.30
3	3*0.30=0.90
4	4*0.35=1.40
5	5*0.10=0.50
	3.20

"On average an oatmeal cookie has 3.2 Raisins"

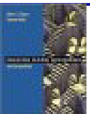

Variance and Standard Deviation

We know the random variable Y has many possible outcomes. If you were forced to give a **"BEST GUESS" for the uncertainty in Y**, what number would you give?

Some people prefer to give the **range of the outcomes of Y**, i.e. the MAX VALUE of Y minus the MIN VALUE of Y.

However, this completely ignores that some values of Y may be more likely than others.

SUGGESTION:

Calculate the "BEST GUESS" for the DISTANCE from the MEAN. **The standard deviation** can be thought of such a guess. The standard deviation of Y is **the square root of the variance** of Y.

Variance and Standard Deviation

Variance:

$$Var(Y) = \sigma_Y^2 = E[(Y - E[Y])^2]$$
$$= E[(Y^2 - 2 \cdot Y \cdot E[Y] + E^2[Y])]$$
$$= E[Y^2] - 2 \cdot E[Y] \cdot E[Y] + E^2[Y]$$
$$= E[Y^2] - E^2[Y]$$

Standard Deviation:

$$\sigma_Y = \sqrt{\sigma_Y^2} = \sqrt{E[Y^2] - E^2[Y]}$$

Draft: Version 1

Variance and Standard Deviation

- If Z=aY+b, a,b constants, Y a rv:

$$Var(Z) = a^2 Var(Y)$$

- If Z=aX+bY, a,b const., X,Y indepent rv's:

$$Var(Z) = a^2 Var(X) + b^2 Var(Y)$$

Example:

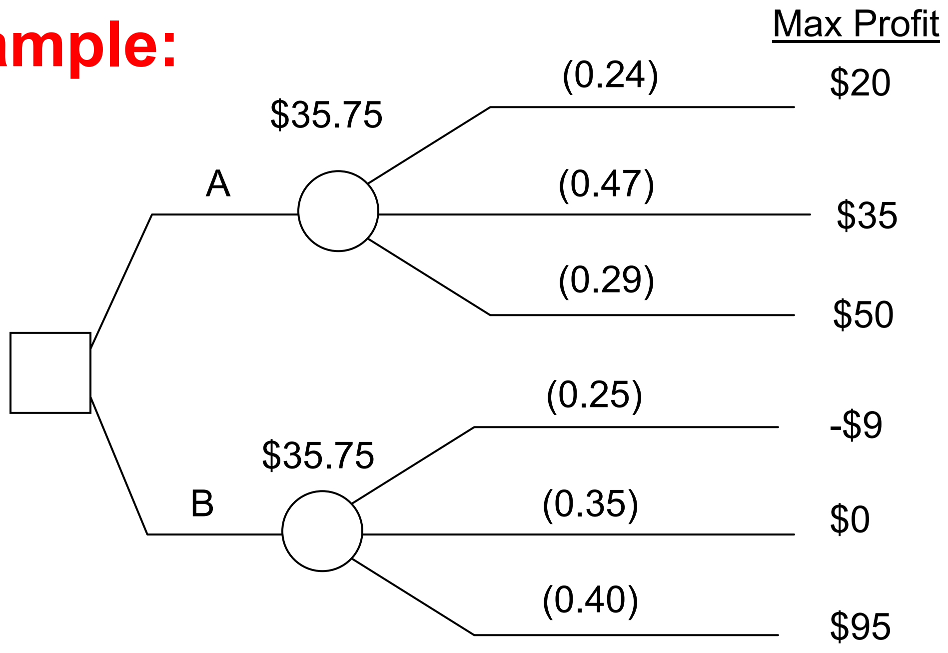

Note that:

$$E[A] = E[B]$$

$$\Pr(\text{Profit} \le 0 \mid A) = 0;$$

$$\Pr(\text{Profit} \le 0 \mid B) = 0.6$$

Draft: Version 1

Variance and Standard Deviation

Alternative A						
Prob	Profit	Profit^2	Prob*Profit	Prob*(Profit^2)	Variance	St. Dev
0.24	20	400	4.80	96.00		
0.47	35	1225	16.45	575.75		
0.29	50	2500	14.50	725.00		
		$E[Y]=$	**35.75**			
			1278.0625	**1396.75**	**118.69**	**10.89438**
			$=E^2[Y]$	$=E[Y^2]$	$=E[Y^2]-E^2[Y]$	σ_Y

Alternative B						
Prob	Profit	Profit^2	Prob*Profit	Prob*(Profit^2)	Variance	St. Dev
0.25	-9	81	-2.25	20.25		
0.35	0	0	0.00	0.00		
0.4	95	9025	38.00	3610.00		
			35.75			
			1278.0625	**3630.25**	**2352.19**	**48.49936**

Draft: Version 1

Example:
Oil Wildcatter
Problem Continued

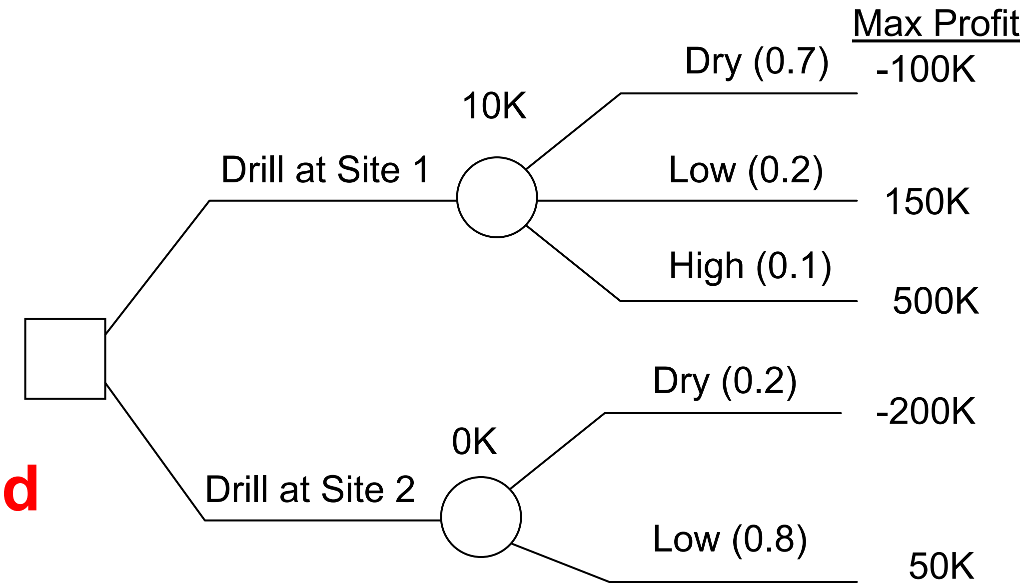

Drill at Site 1						
Prob	**Profit**	**Profit^2**	**Prob*Profit**	**Prob*(Profit^2)**	**Variance**	**St. Dev**
0.7	-100	10000	-70.00	7000.00		
0.2	150	22500	30.00	4500.00		
0.1	500	250000	50.00	25000.00		
			10.00	$= E[Y]$		
			100	**36500.00**	**36400.00**	**190.7878**

$= E^2[Y]$ $\quad$ $= E[Y^2]$ $\quad$ σ_Y

$= E[Y^2] - E^2[Y]$

Draft: Version 1

Expected Value, Variance and Standard Deviation

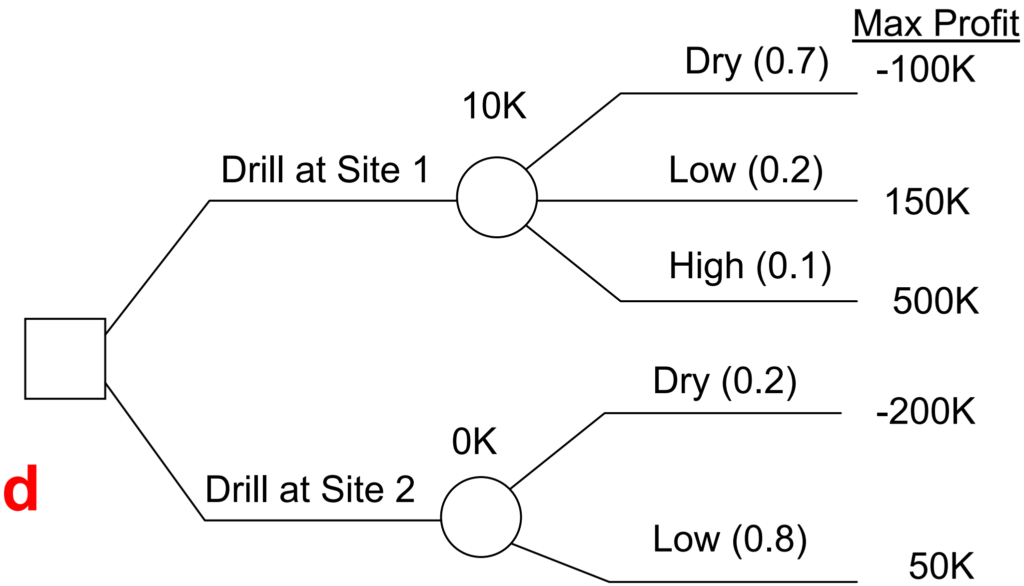

Example: Oil Wildcatter Problem Continued

Drill at Site 2						
Prob	**Profit**	**Profit^2**	**Prob*Profit**	**Prob*(Profit^2)**	**Variance**	**St. Dev**
0.2	-200	40000	-40.00	8000.00		
0.8	50	2500	40.00	2000.00		
			0.00	$\rightarrow = E[Y]$		
			0	**10000.00**	**10000.00**	**100**

$= E^2[Y]$ $\quad = E[Y^2]$ $\quad \sigma_Y$

$= E[Y^2] - E^2[Y]$

Draft: Version 1

Expected Value, Variance and Standard Deviation

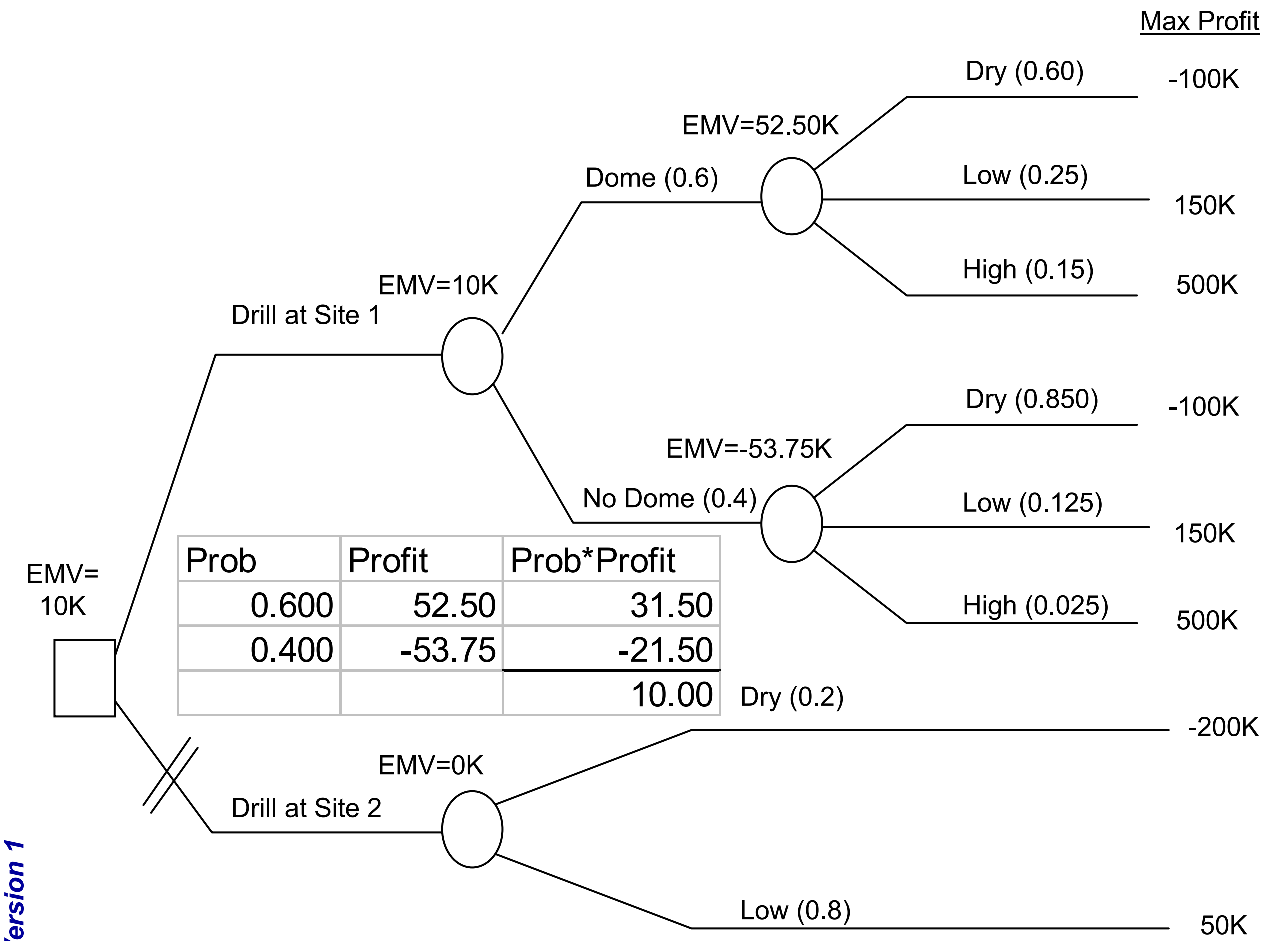

Prob	Profit	Prob*Profit
0.600	-100.00	-60.00
0.250	150.00	37.50
0.150	500.00	75.00
		52.50

Prob	Profit	Prob*Profit
0.850	-100.00	-85.00
0.125	150.00	18.75
0.025	500.00	12.50
		-53.75

Prob	Profit	Prob*Profit
0.600	52.50	31.50
0.400	-53.75	-21.50
		10.00

Prob	Profit	Prob*Profit
0.200	-200.00	-40.00
0.800	50.00	40.00
		0.00

Draft: Version 1

Expected Values can be calculated by **"folding back the tree"**

Let: X = The failure time of a component

Definition: Let X be a continuous rv. Then a probability density function (pdf) of X is a function f(x) such that for any two numbers a and b with a < b:

$$\Pr(X \in [a,b]) = \int_a^b f(x)dx$$

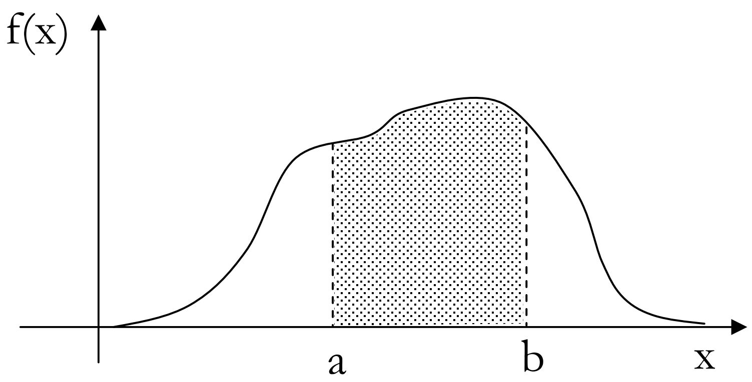

For f(x) to be a legitamate pdf we must have:

$$f(x) \geq x, \text{ for all possible values } x$$

$$\int_{-\infty}^{\infty} f(x)dx = 1$$

Cumulative distribution function:

$$F(x) = \Pr(X \le x) = \int_{-\infty}^{x} f(u)du$$

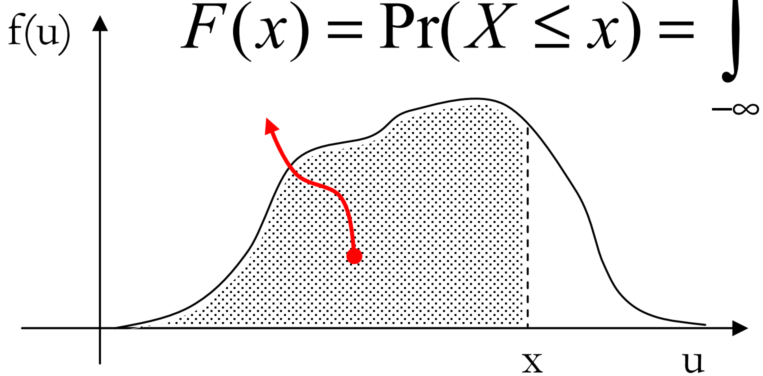

The p-th quantile x_p:

$$F(x_p) = \Pr(X \le x_p) = p$$

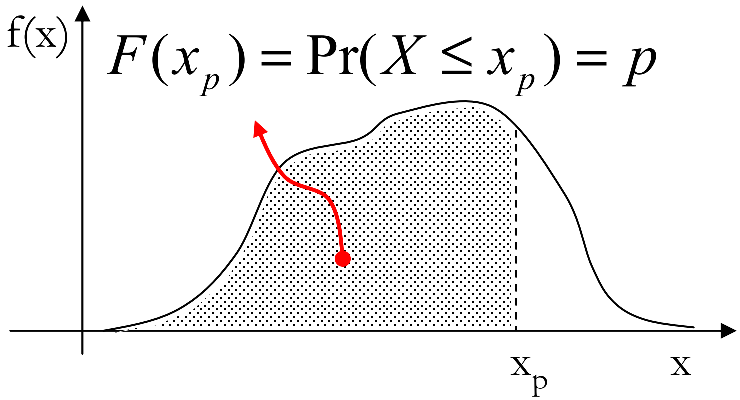

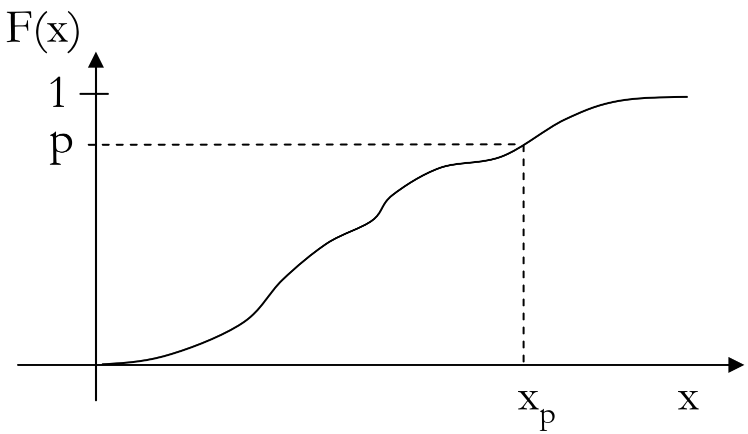

Draft: Version 1

Continuous Probability Distributions

Let: X = The failure time of a component

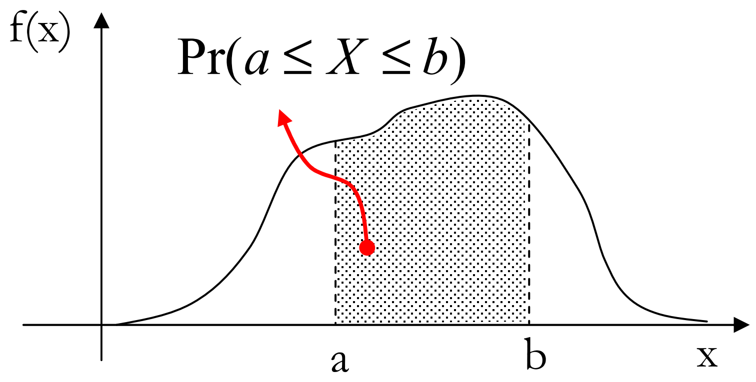

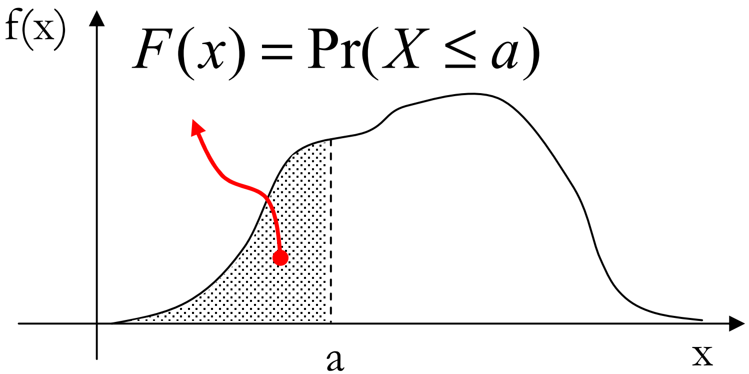

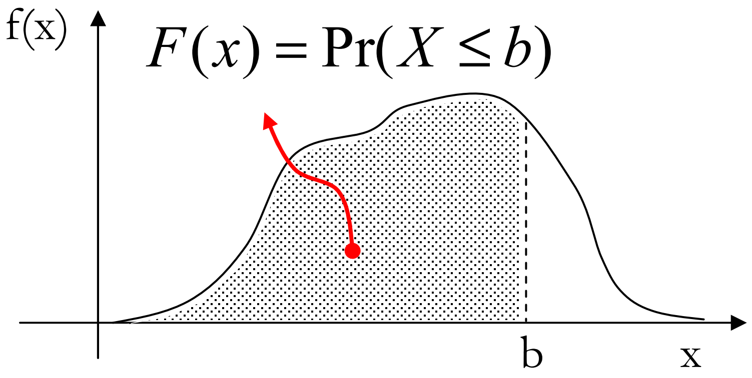

Conclusion:

$$\Pr(a \le X \le b) = \Pr(X \le b) - \Pr(X \le a) = F(b) - F(a)$$

Draft: Version 1

Continuous Probability Distributions

Examples theoretical density functions:

Normal: $$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \cdot \exp\left(-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right), x \in \mathbb{R}$$

Exponential: $$f(x) = \lambda \cdot \exp(-\lambda \cdot x), x > 0$$

Beta: $$f(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\cdot\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}, x \in [0,1]$$

More in Chapter 9

Expected Value: $$E[X] = \int_{-\infty}^{\infty} x f(x) dx$$

Variance: $$Var(X) = \int_{-\infty}^{\infty} (x - E[X])^2 f(x) dx$$

Formulas carry over from the discrete to the continuous case

Draft: Version 1

Dominance Revisited

DETERMINISTIC DOMINANCE

Assume random Variable X Uniformly Distributed on [A,B]

Assume random Variable Y Uniformly Distributed on [C,D]

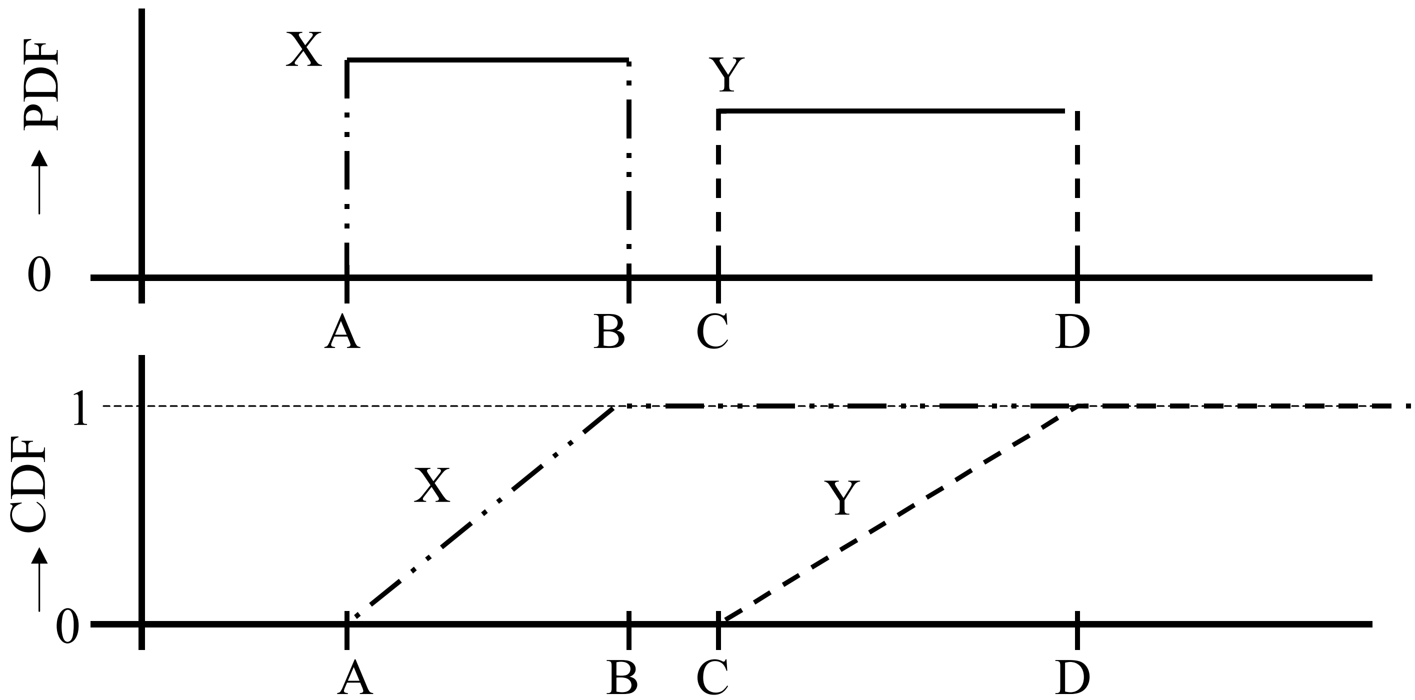

Dominance Revisited

STOCHASTIC DOMINANCE

Assume random Variable X Uniformly Distributed on [A,B]

Assume random Variable Y Uniformly Distributed on [C,D]

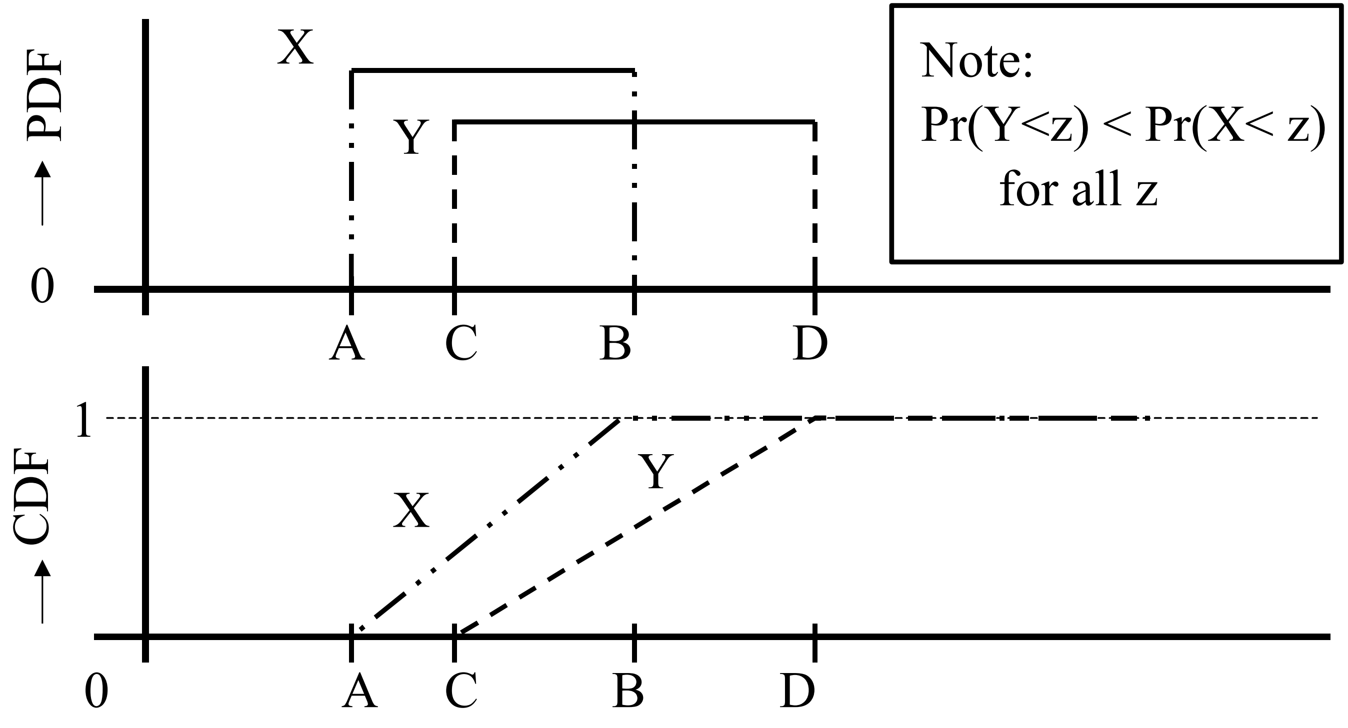

Note:
Pr(Y<z) < Pr(X< z)
for all z

Dominance Revisited

CHOOSE ALTERNATIVE WITH BEST EMV

Assume random Variable X Uniformly Distributed on [A,B]

Assume random Variable Y Uniformly Distributed on [C,D]

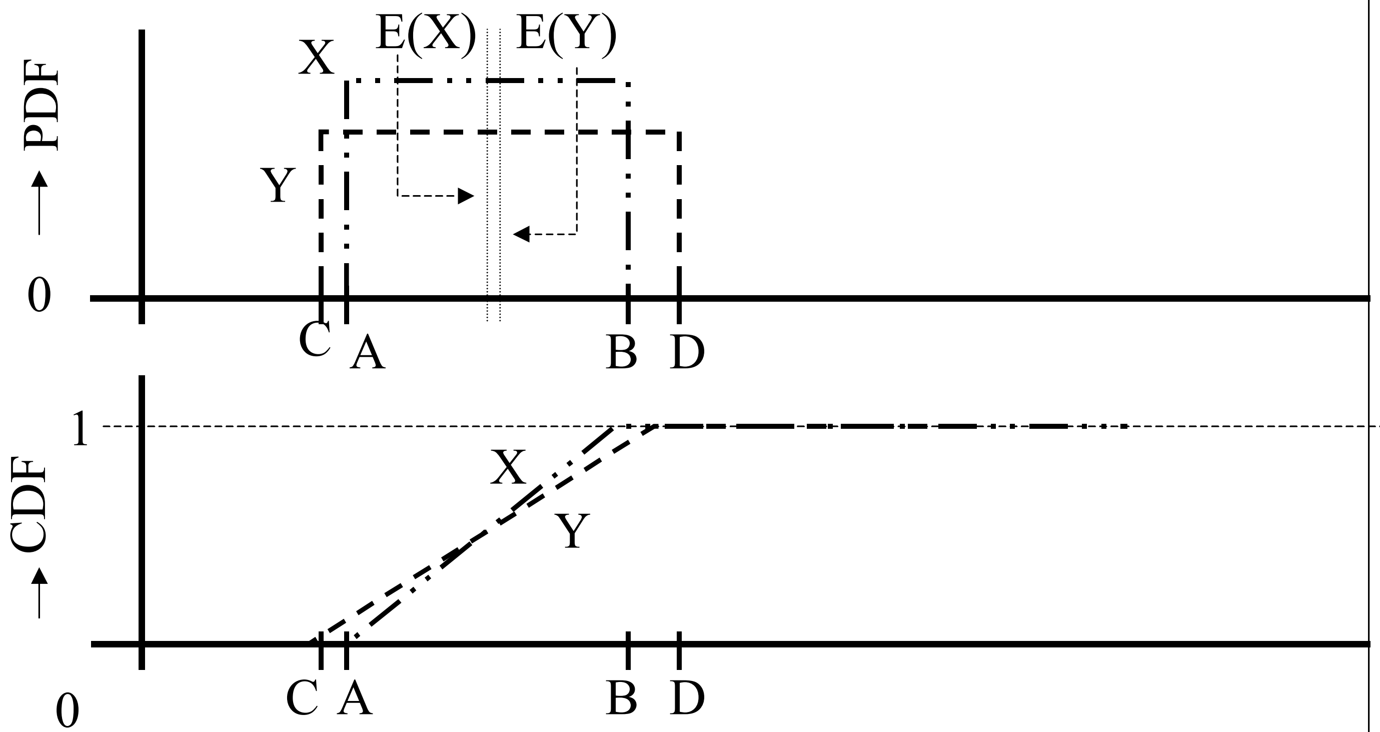

Making Decisions under Uncertainty

Deterministic Dominance Present

↓

Stochastic Dominance Present

↓

Making Decisions based on EMV

Chances of **an unlucky outcome** increase

Draft: Version 1