Prof. Massimiliano Ferrara

Principles of Business Analytics

To Enter or not to Enter?

- **Book Browser**, Inc. is evaluating whether or not to commit in a new division which specializes in the **web-based sale of rare books**.
- They have already spent money in the investigation of this opportunity, but they have to spend additional **2MUSD** if they decide to enter the market.
- They believe that there is a **30% chance** that no competitors will enter the market.
- In that case, they expect to earn **$15 millions**
- If one competitor enters (which they believe happens with P=0.60), BB expects to earn **6MUSD**
- BB expects that **two *competitors*** will enter the market with 0.1 probability.
- In that case, BB expects **losses for 2MUSD**
- If BB does not enter the market, it can have sure profits from other initiatives equal to **5MUSD**

Decision Problem Elements

- ***OBJECTIVES***
- ***ATTRIBUTES***
- ***DECISION ALTERNATIVES***
- ***EVENTS***
- ***CONSEQUENCES***

The Elements

- ***OBJECTIVES:***
  - Maximize Expected Profit
- ***ATTRIBUTES:***
  - Money
- ***ALTERNATIVES:***
  - Enter (Risky)
  - Do not Enter (Sure)
- ***EVENTS:***
  - N° of competitors
- ***CONSEQUENCES:***
  - Profit or Loss

The Payoff Table for Book Browser

	0	1	2 o +
A1: Enters	15-2=13	6-2=4	-2-2=-4
A2: Does not Enter	5	5	5
Probability	0,3	0,6	0,1

Bocconi

Influence Diagrams and Decision Trees

Influence Diagrams

- **Influence diagrams (IDs) are…**

  "a graphical representation of decisions and uncertain quantities that explicitly reveals probabilistic dependence and the flow of information"

- **ID formal definition:**
  – ID = a network consisting of a directed graph G=(N,A) and associated node sets and functions (*Schachter, 1986*)

Elements of an ID

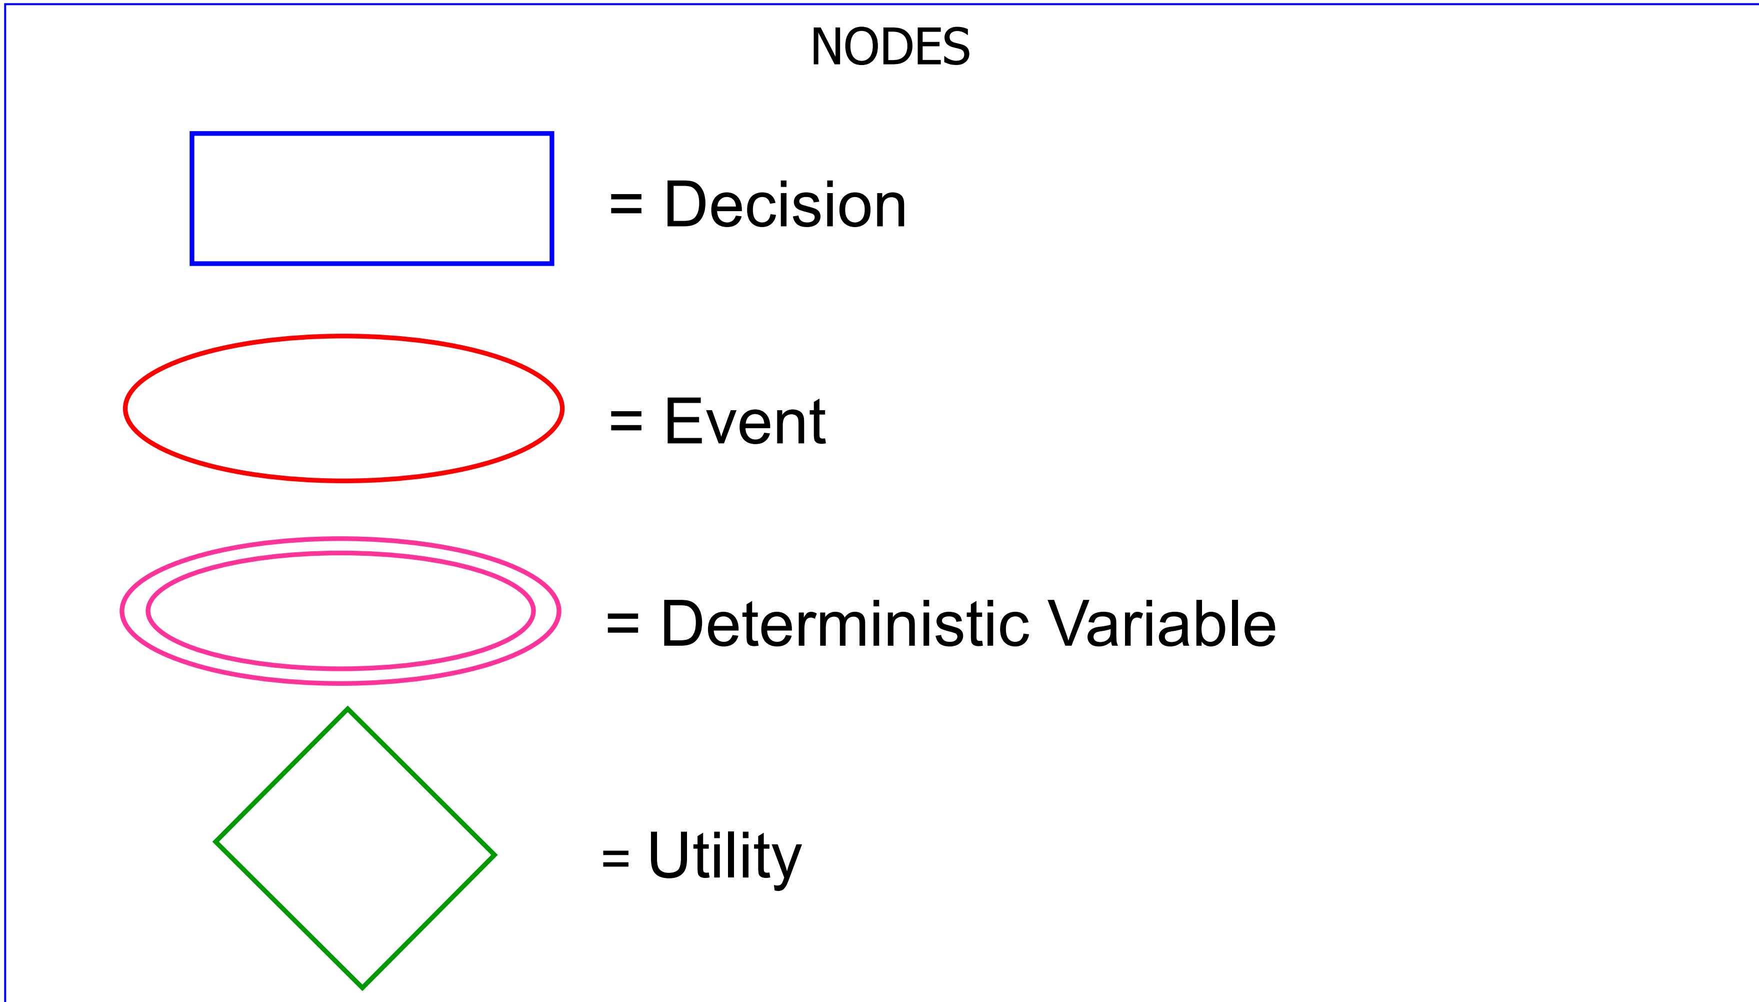

Arcs

- Connect two nodes (initial and final)
- Reveal a dependence between the nodes

- **Informational**

- Ends in a Decision Node
- Reveals that the decision-maker is aware of the corresponding information when taking the decision

- **Functional**

- Ends in the Utility Node
- Signal that the final result depends on the outcomes of the preceding node(s)

Arcs (cont.)

Conditional (probabilistic dependence)

- Ends in an Event node
- Reveals that the random quantity in the final node depends probabilistically from the argument of the initial node

OR

- Ends in a deterministic node
- Signals that the quantity in the final node depends in a deterministic fashion from the argument (eventually random) of the initial node

A generic ID

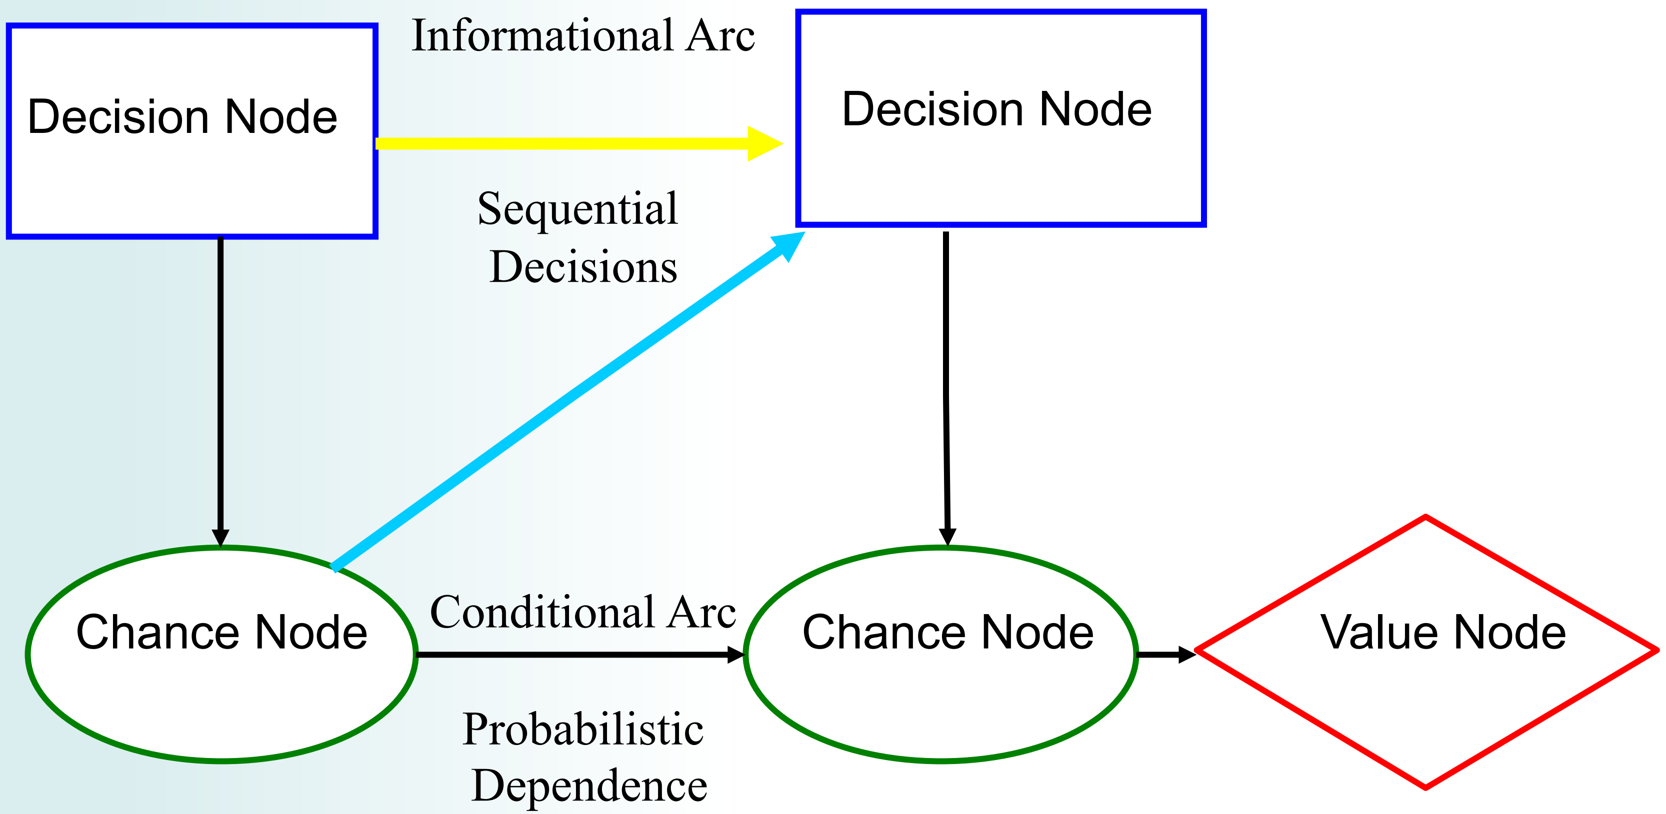

Decision Node
Informational Arc
Decision Node
Sequential
Decisions
Chance Node
Conditional Arc
Chance Node
Value Node
Probabilistic
Dependence

ID for Book Browser

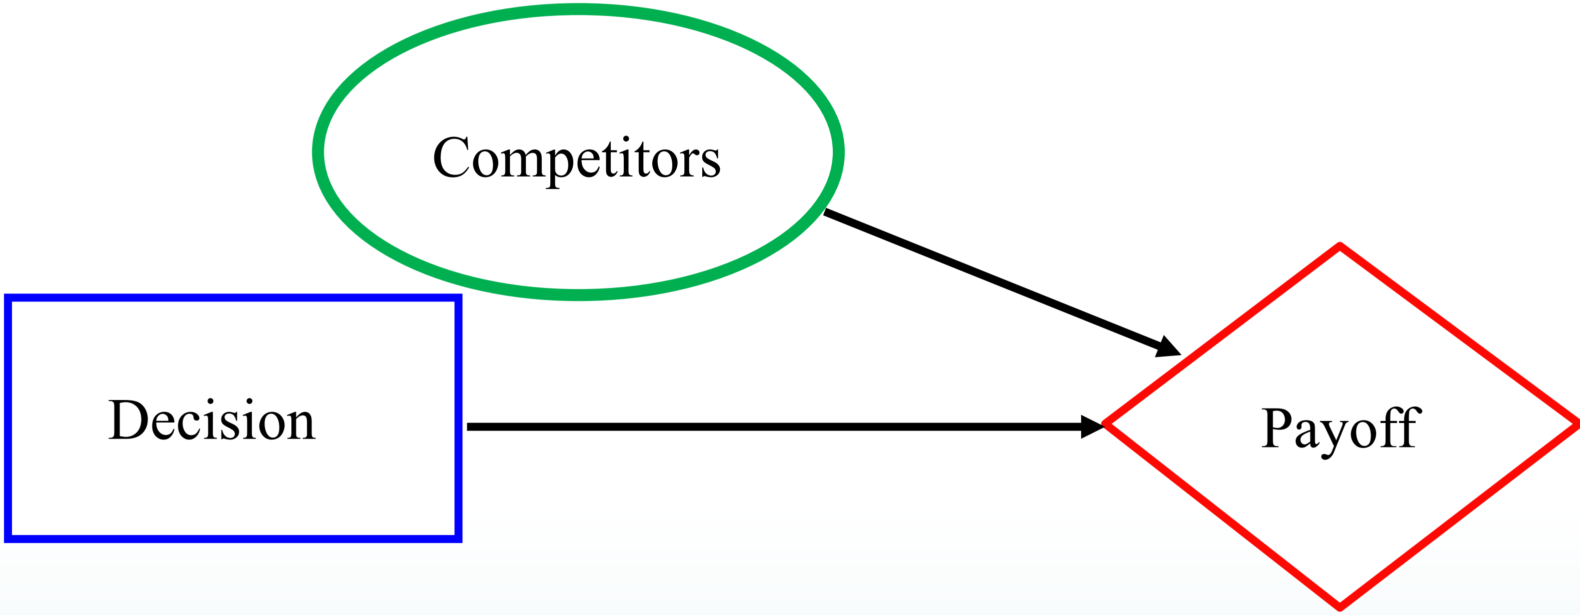

Decision Trees

Bocconi

- **Decision Trees** (DTs) are made of by the same type of arcs of Influence Diagrams, but **highlight** all the possible **event combinations**.
- Instead of arks, one finds **branches** that emanate from the nodes as many as the **Alternatives or Outcomes of each node.**
- With respect to Influence Diagrams, Decision Trees have the advantage of **showing all possible patterns**, but their **structure** becomes quite **complicated** at the **growing** of the **problem complexity.**

The DT for BB

The corresponding DT

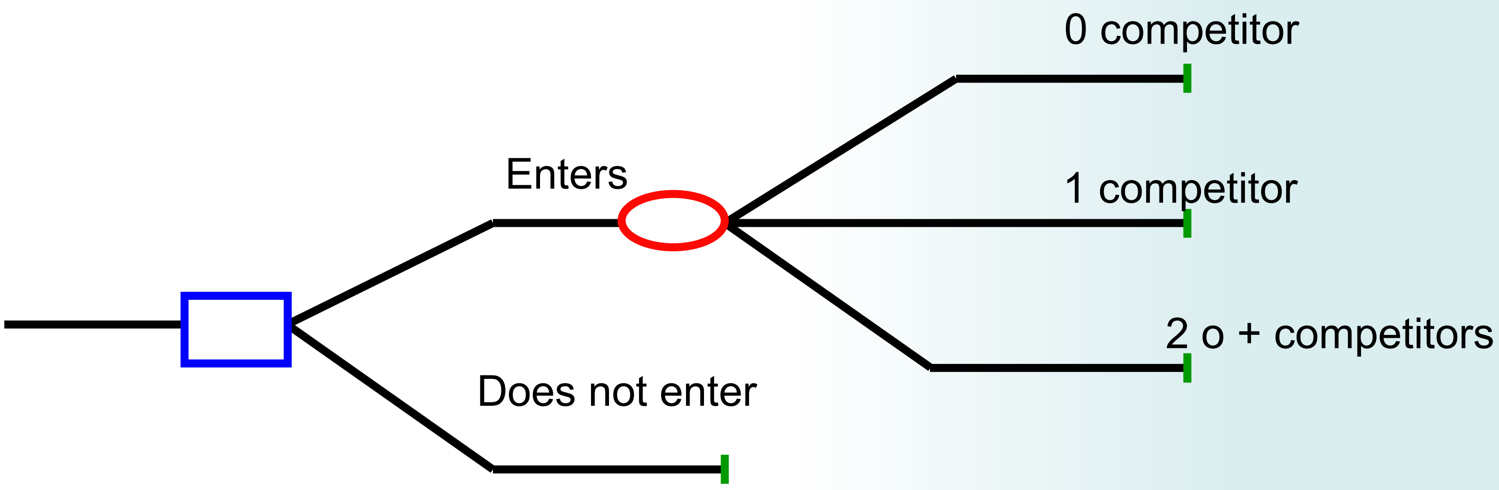

Let us insert the numbers

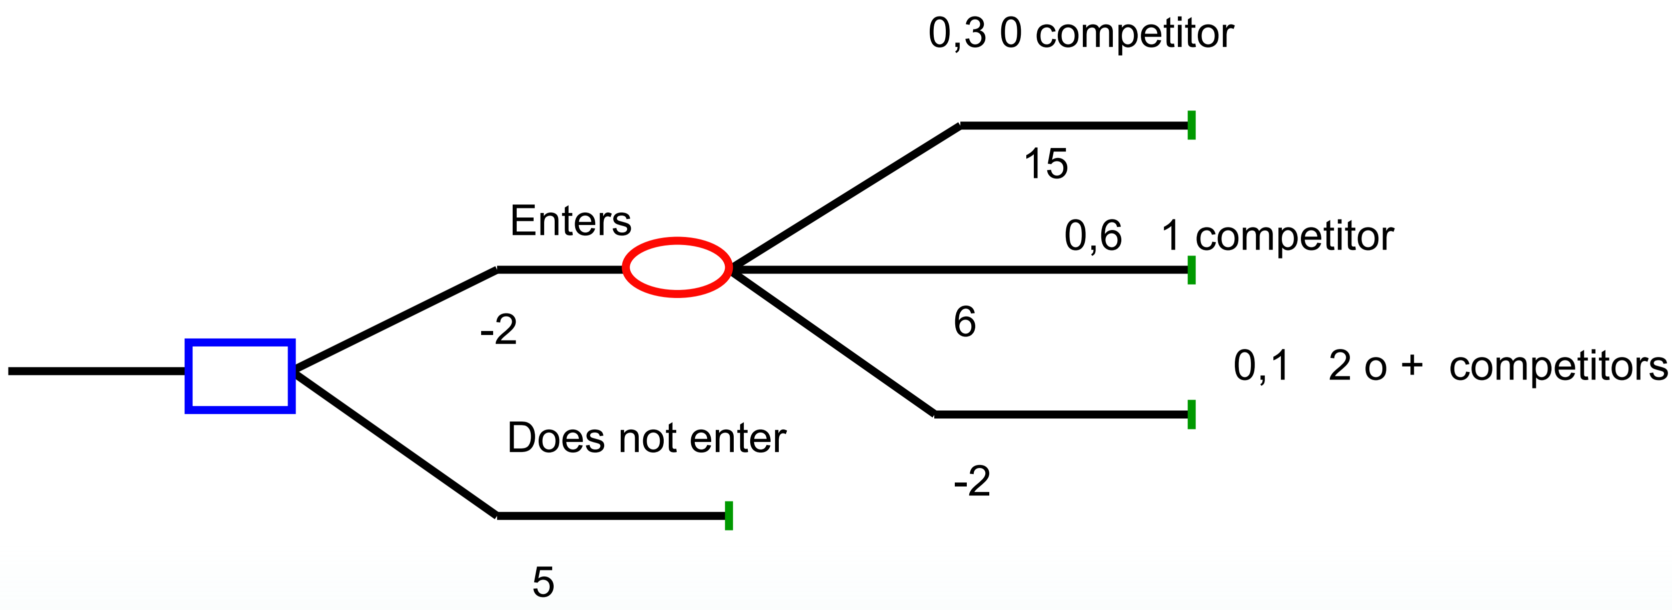

- The data are reported directly in the tree (in black)

Elements of a Decision Tree

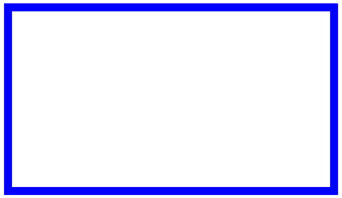

Decision Node: the branches explicitly display the available alternatives

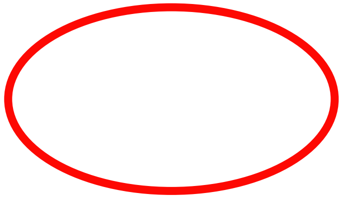

Chance Node: the branches display the possible outcomes and their probabilities

Terminal Node: in Tree plan it is used this shape, other software may use a different shape

Floding-back the tree

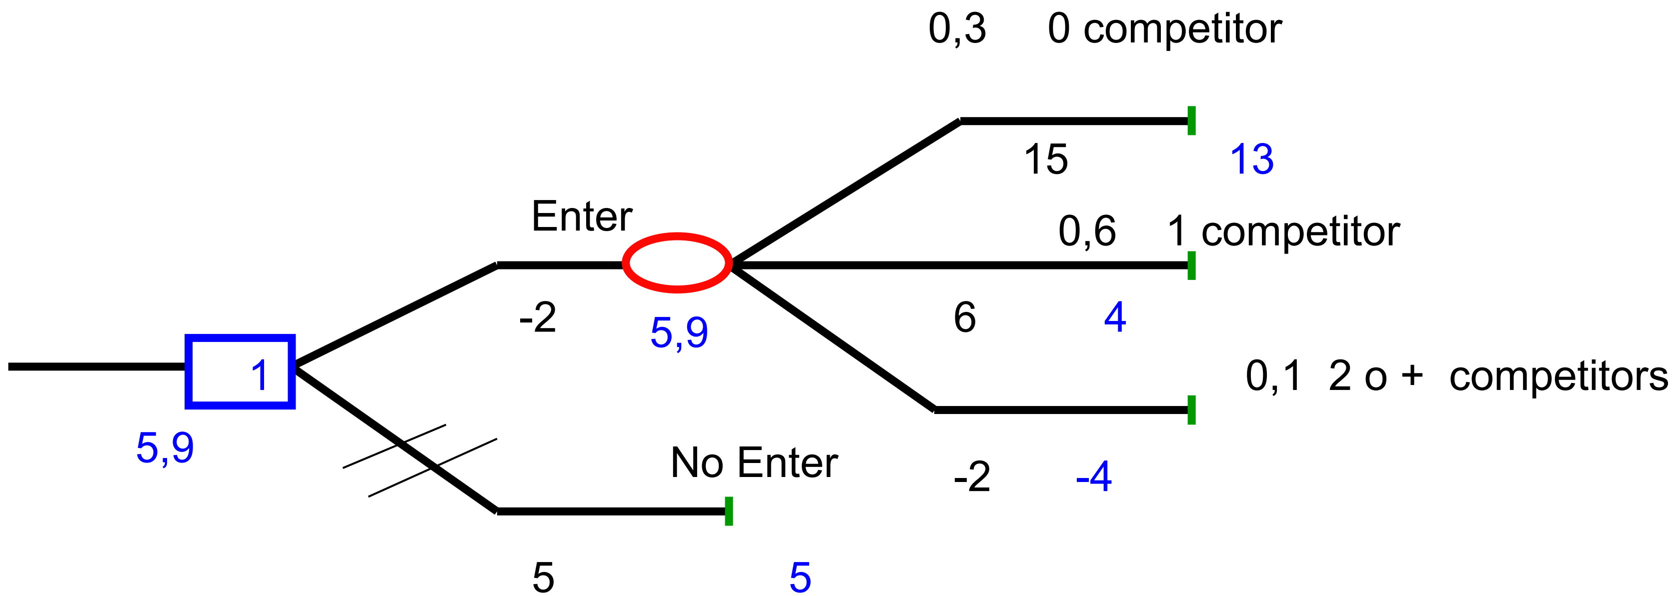

We illustrate the calculations on the board.
The preferred alternative is 1

Book Browser continued

BB objective: **maximize**

net expected profit

E[A1]=0,3X13+0,6x4+0,1x(-4)=5,9

E[A2]=5

Preferred Alternative:

enter the market Maximum expected profit: 5,9MUSD

Comments on DTs

- They reveal all possible combinations **of events, outcomes and decisions**
- As IDs, they are **directed graphs**, drawn and read from the left to the right
- Their **complication**, however, **increases** exponentially with the number of nodes
- They allow also to include the **effect of information**

Bocconi

Università Commerciale Luigi Bocconi

The GM Problem

- Game Magic, Inc. (GM) is developing a new platform for a computer game play station.
- Market demand for the new platform is uncertain.
- Problem Data

They are concerned about when they should enter the market:

1

They could continue the development of the platform for one more quarter.

Could enter the market this quarter and be guaranteed to have "first-mover" advantages. (A-Starry will *not* enter this quarter.)

2

GM Problem DATA

	A	B	C	D	E
1	**Game Magic, Inc.**				
2					
3	**Introduce Now**				
4					
5	**Demand**	**Prob**	**Profit**		
6	High	0.2	50		
7	Low	0.8	-10		
8					
9	**Introduce Next Quarter**				
10		A-Starry Enters		A-Starry Does Not Enter	
11	**Demand**	**Prob**	**Profit**	**Prob**	**Profit**
12	High	0.4	40	0.7	120
13	Low	0.6	-100	0.3	5
14					
15	**Prob(Competitor Enters)**		0.35		
16	**Development cost**		-2		
17					

ID for GM

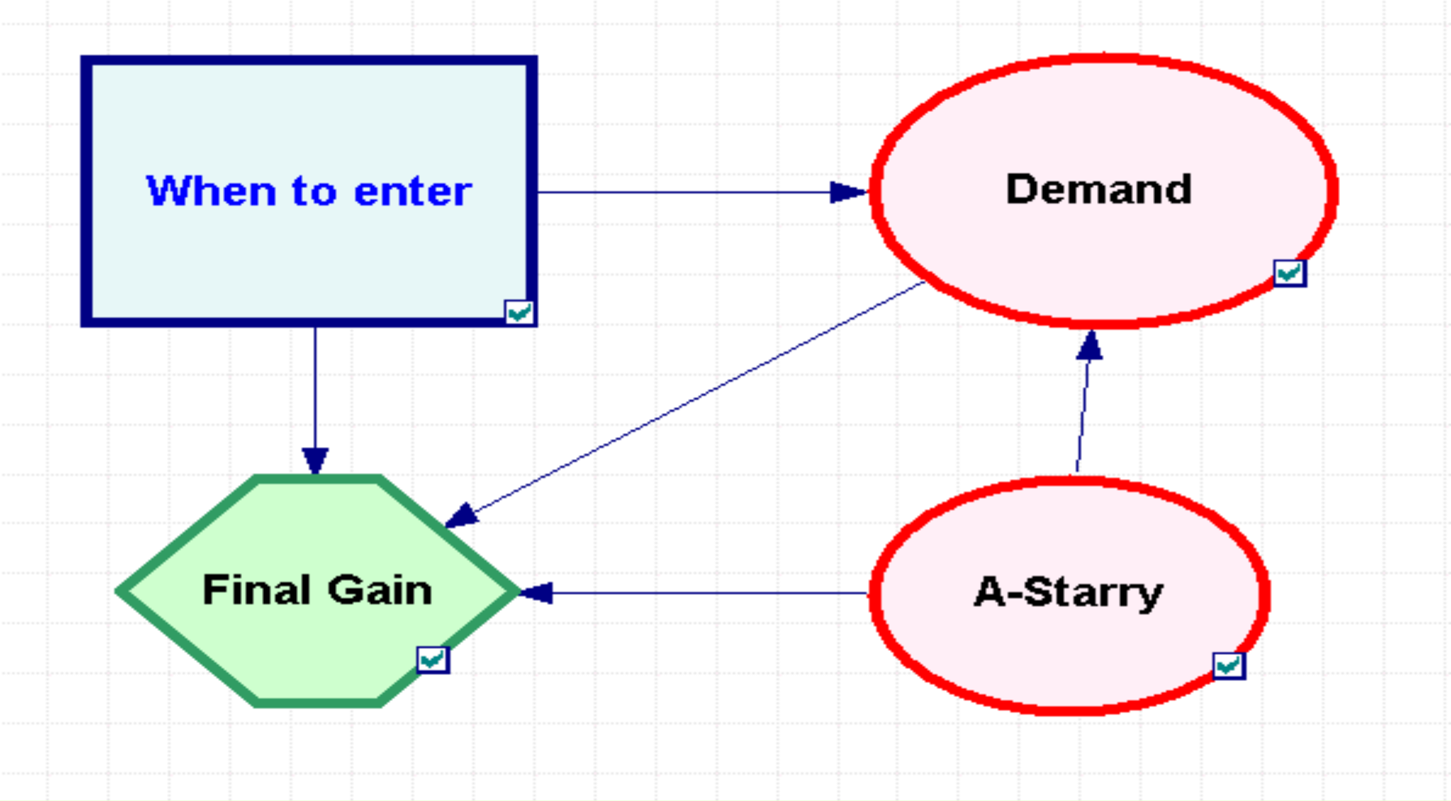

Decision Tree for GM

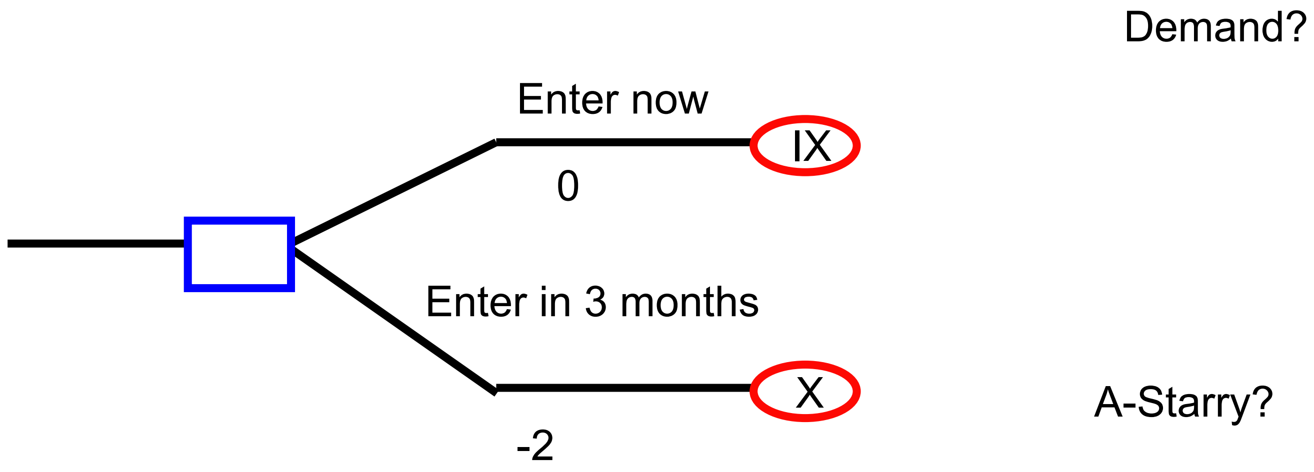

Developing the Decision Tree

A1: *GM enters the market now*

Chance node IX

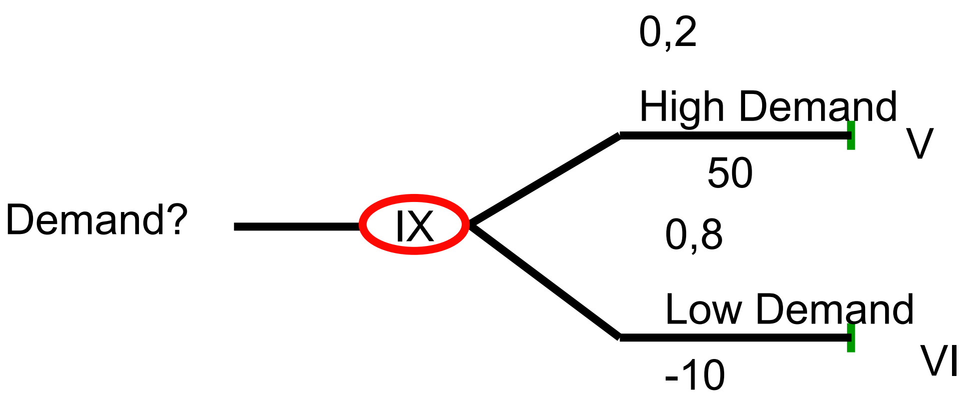

Developing the DT for GM

A2: *GM enters in a quarter*

↓

Chance node X

Chance Nodes VII e VIII

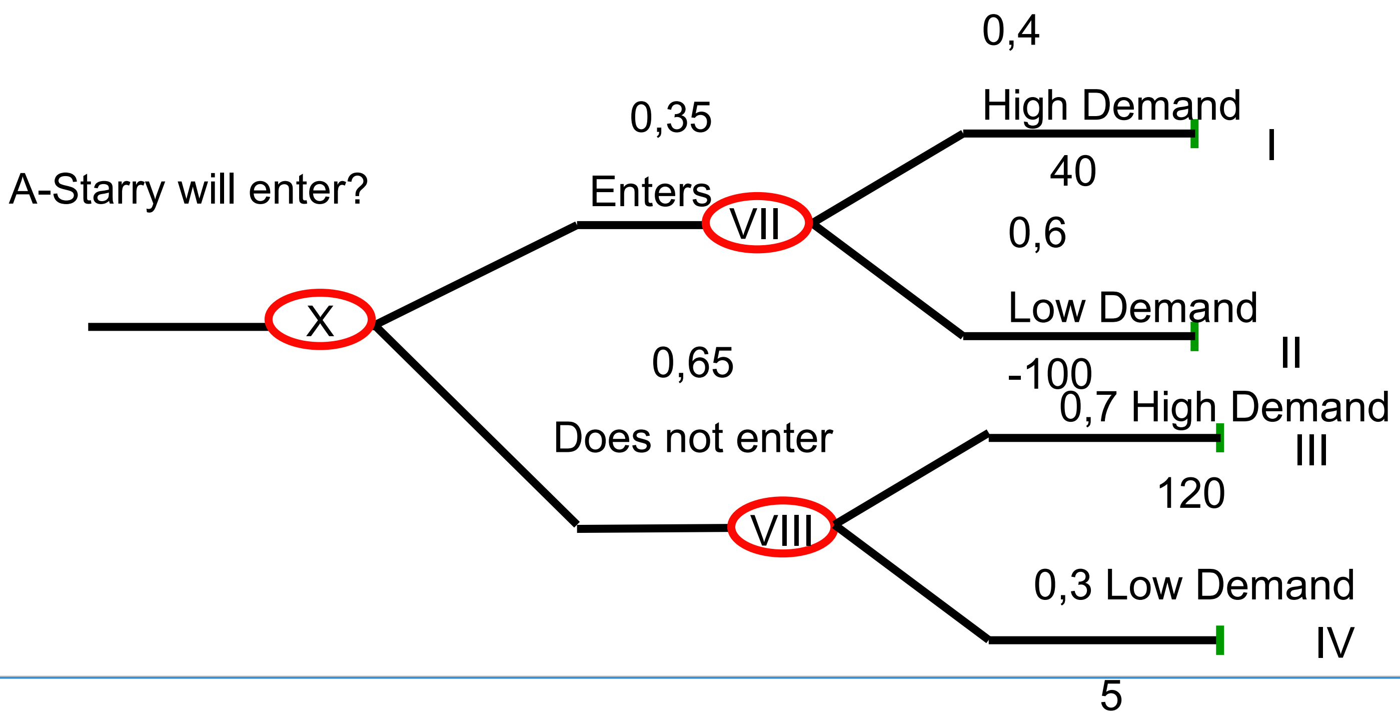

The Complete Tree

GM

Conditional Probability (1)

- *Note what happens to the probabilities of nodes VII and VIII.*

We can define the following outcomes:

- *E*: "A-Starry enters"
- *D*: "A-Starry delays entering"
- *H* evento "High Demand"
- *L* evento "Low Demand".

Conditional Probability

- The probability of high demand depends on whether **A-Starry has enters or not the market.**
- More precisely, the fact that A-Starry enters the market influences **Your view on the probability of High Demand**.
- In other words:

  P(High Demand | A-Starry)

Conditional Probability: the probability of event A <u>given</u> that event B has happened

Conditional Probability

- Suppose the decision process is at **Node VII**. We *know* **A-Starry has entered the market** and use this information to revised our estimates of **market demand**.
- Given datum: $P(H|E) = 0.40$
- $P(L|E) = 1 - P(H|E) = 1 - 0.40 = 0.60$.
- Given datum: $P(H|D) = 0.70$
- So that: $P(L|D) = 1 - P(H|D) = 1 - 0.70 = 0.30$.
- **Finally, at Node X**, $P(E) = 0.35$ and $P(D) = 1 - 0.35 = 0.65$.

Folding-back (Roll-back) the DT

- Recall that the objective is **maximize total expected profit.**
- Folding back a decision tree means **assigning values** to each of the nodes in the tree.

Value of a Node (When Maximizing Expected Profit):

- The value that is assigned to any node represents the **maximum expected profit** that can be generated from this point on in the decision process.

Value Assignment Rules

Node Type	Rule for Assign a Value to the Node
Terminal Node	The value of each terminal node is sum of the rewards and costs that are on the unique path that connects the first node of the tree to this terminal node.
Event Node	For each arc emanating from an event node, multiple the value of the node at the other end of the arc by the probability associated with the arc and add all such products over all arcs emanating from the node.
Decision Node	The value of the decision node is the maximum of the values at the end of each arc emanating from this node.

Identifying an Optimal Strategy

- At each decision node, take the branch associated with the highest value at the other end of the branch. This highest value is, of course, the value of this decision node.
- Sometimes useful to put **slashes** through those arcs emanating from a decision node that should not be traversed.
- In our final decision tree, we see that the **optimal strategy is:**

Develop the product further. Introduce the product next quarter.

Sensitivity Analysis

- Sensitivity Analysis is an **essential tool** to derive managerial insights from models
- It allows to fully **exploit the modelling efforts**
- In this class:
- **One way** sensitivity analysis **on one** parameter
- **Two way** sensitivity analysis **on two** parameters

One Way Sensitivity Analysis

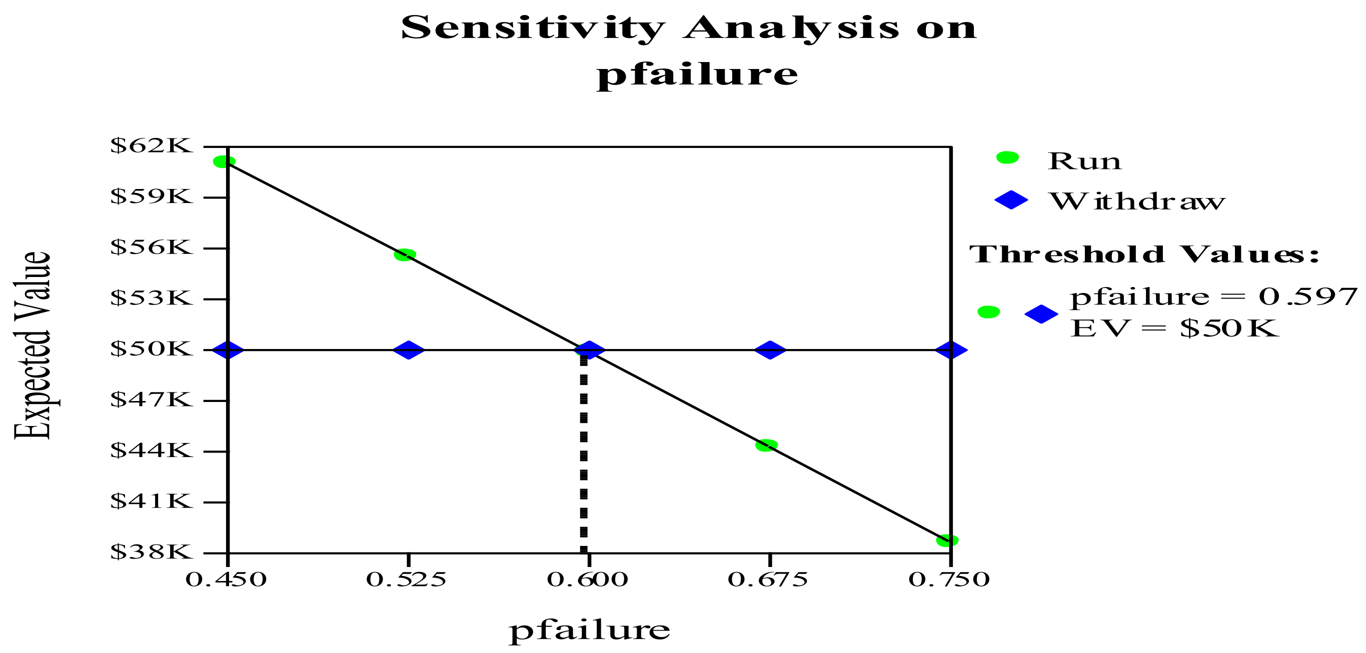

1) Idenfity the parameter of interest
2) Keep all parameters fixed but the one of interest
3) Plot the expected value of each alternative as a function of the parameter of interest

The threshold value is the value of the parameter at Which two alternatives become equally preferred.

Insights from One-Way Sensitivity Plot

- ***INDIFFERENCE:*** we find the value of the parameter (probability) that makes us indifferent among the alternatives
- ***STABILITY:*** we can understand if we are close or not to indifference and therefore if a small change in the inputs can make use change our mind
- ***TREND:*** we understand whether the expected payoff of the alternative is increasing or decreasing as a function of the parameter under scrutiny

Tornado Diagrams

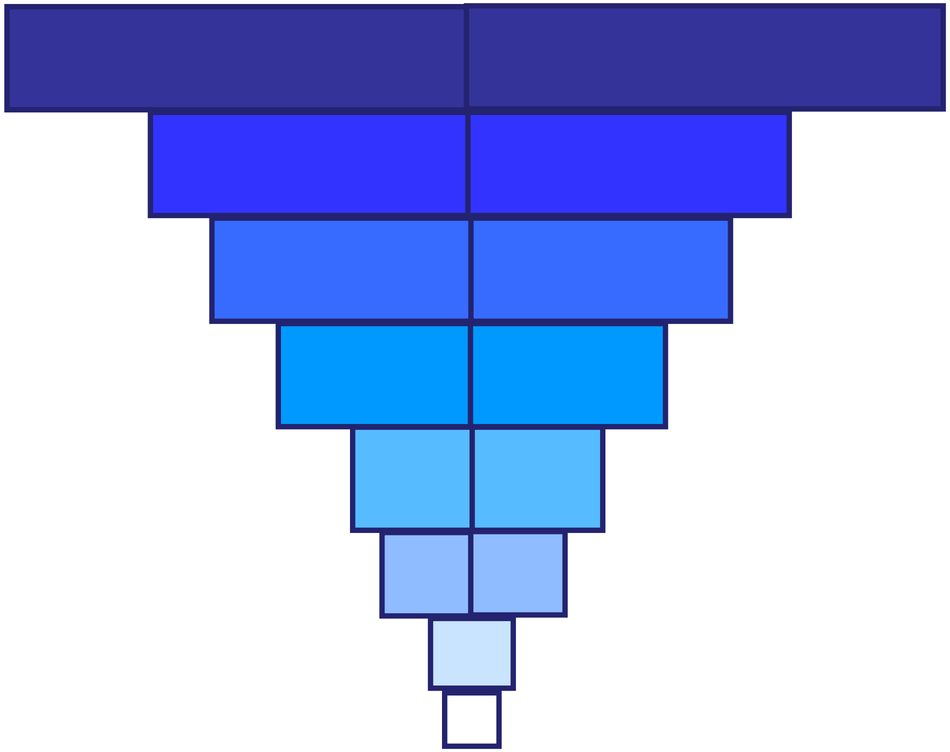

Tornado Diagrams

- Consider that the **manager is asking**:
  - What if a given probability increases by 20%? And
  - What if revenues increase by 35%?
- What is the **impact** of these variations on the **expected profit**?
- To answer these questions the **business community** relies on a nice graphical representation called a **Tornado Diagram**

How to build a Tornado Diagram

- Select the **variables of interest**
- (if you want to find the most important, you should **select them all**)
- Assign a **variation interval**.

For instance, in GM we wish to see what happens if the probability that A-Starry enters changes from 0.35 to 0.42 and from 0.35 to 0. 28. That is, we consider two variations, one "at the right" and one "at the left". The variation interval is [0.28,0.35] and we evaluate the model at the extremes

We also wish to see what happens when the profit when demand is high and A-Starry enters (we decide to introduce later) (currently at the value of 40) varies from 50 to 60 and from 50 to 40. I.e., we assign a range to this variable [32,48]

We also consider the probability that demand is high given that A-Starry does not enter, varying it of plus or minus 20%

Building Tornado Diagrams (2)

- Then, we need to **evaluate the model** in correspondence of the **variations** (one at a time) of the variables to the **extremes** of their ranges
- In this case, we need **6 evaluations**
- **But variations of what?**

The Choice of the Output

- **The optimal profit**
  - In this case, your output is the maximum value of the problem, irrespectively of what is the preferred alternative
- **The value of the preferred alternative**
  - In this case, you fix the preferred alternative and consider what is the change in expected value given that the variables change
- **Both choices are valid, it depends on the problem**

Table Form

Variable	20% Decrease	20% Increase
Profit A-Starry Ent. (40)	-1.12	1.12
ProbAstarryEnters (0.35)	9.065	-9.065
Prob. High Dem. Given A-Starry Enters (0.7)	-10.465	10.465

Bocconi

Graphical Form

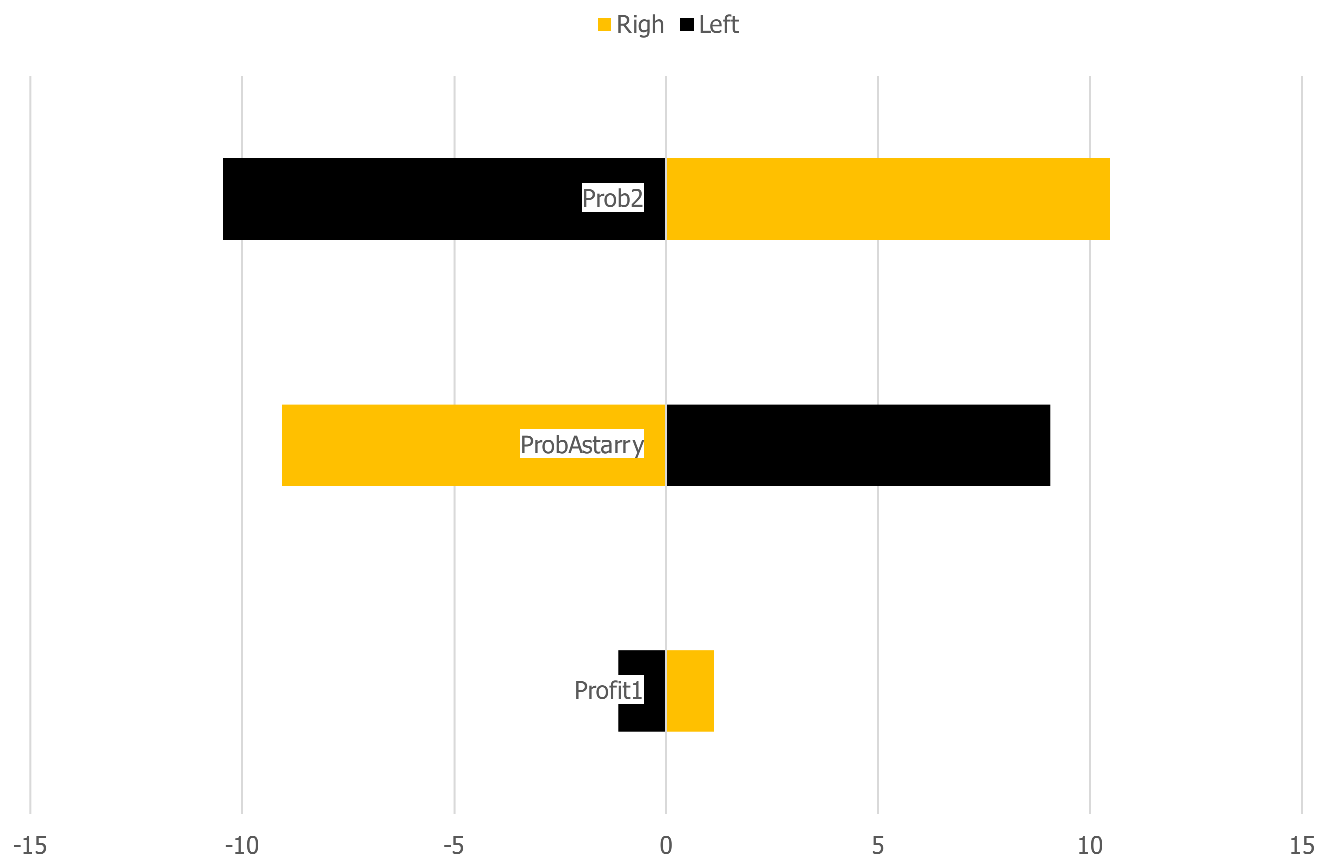

Insights from Tornado Diagrams

- **Direction of impact:** is the change of interest going to increase or decrease my expected profit?
- **Which variable has the greatest impact?**

  Note: if the variations are not all in the same percentage, the insight on ranking becomes shacky
- **Is the preferred alternative changing?**

  This insight is gained, of course, if you select as output the maximum of the decision problem

The value(s) of Information

BB Revised

- **Book Browser, Inc.** is evaluating whether or not they should commit to start a **new division** which specializes in the **web-based** sale of **rare books**.
- They have already spent money in the **investigation** of this opportunity, but they have to spend **additional 2MUSD** if they decide to enter the market.
- They believe that there is a **30%** chance that **no competitors** will enter the market.
- In that case, they expect **to earn $15 milions**
- If one ore more **competitors enter** (which they believe happens with **P=0.70**), BB expects to **earn 5.4MUSD**
- If BB does **not enter** the market, it can have sure profits from **other initiatives** equal to 5MUSD

ID and DT

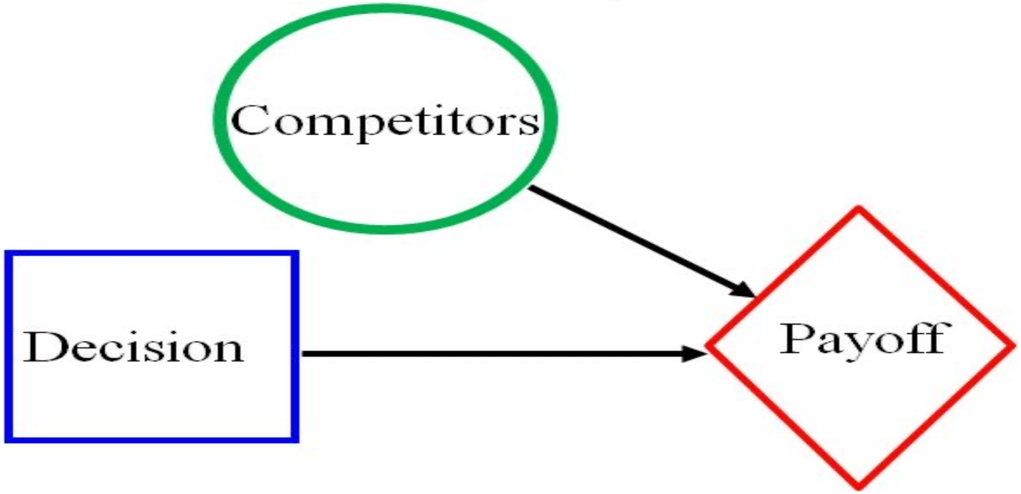
Competitors
Decision
Payoff

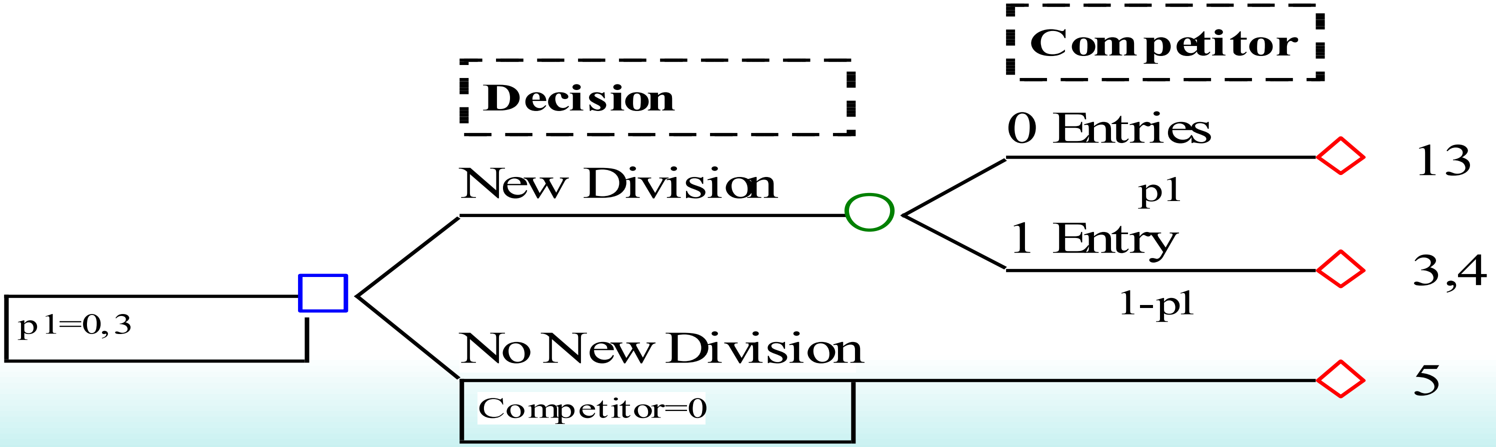
Decision
Competitor
p1=0,3
New Division
0 Entries
p1
13
1 Entry
1-p1
3,4
No New Division
Competitor=0
5

Expected Profit

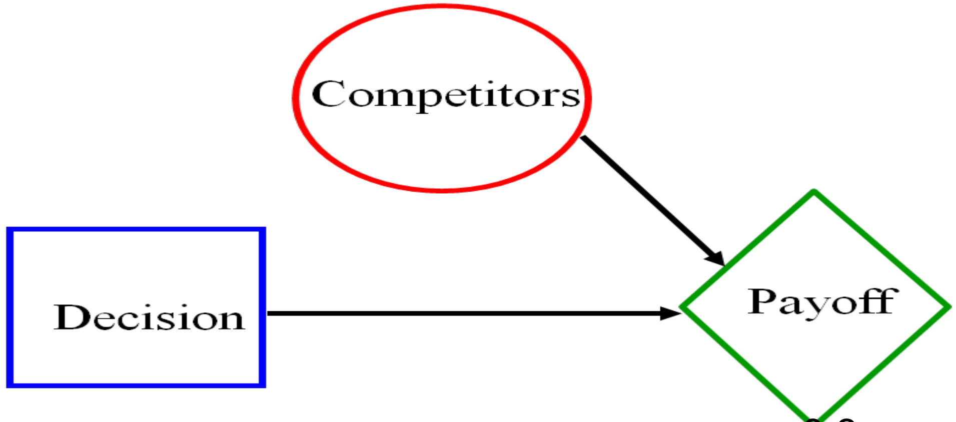

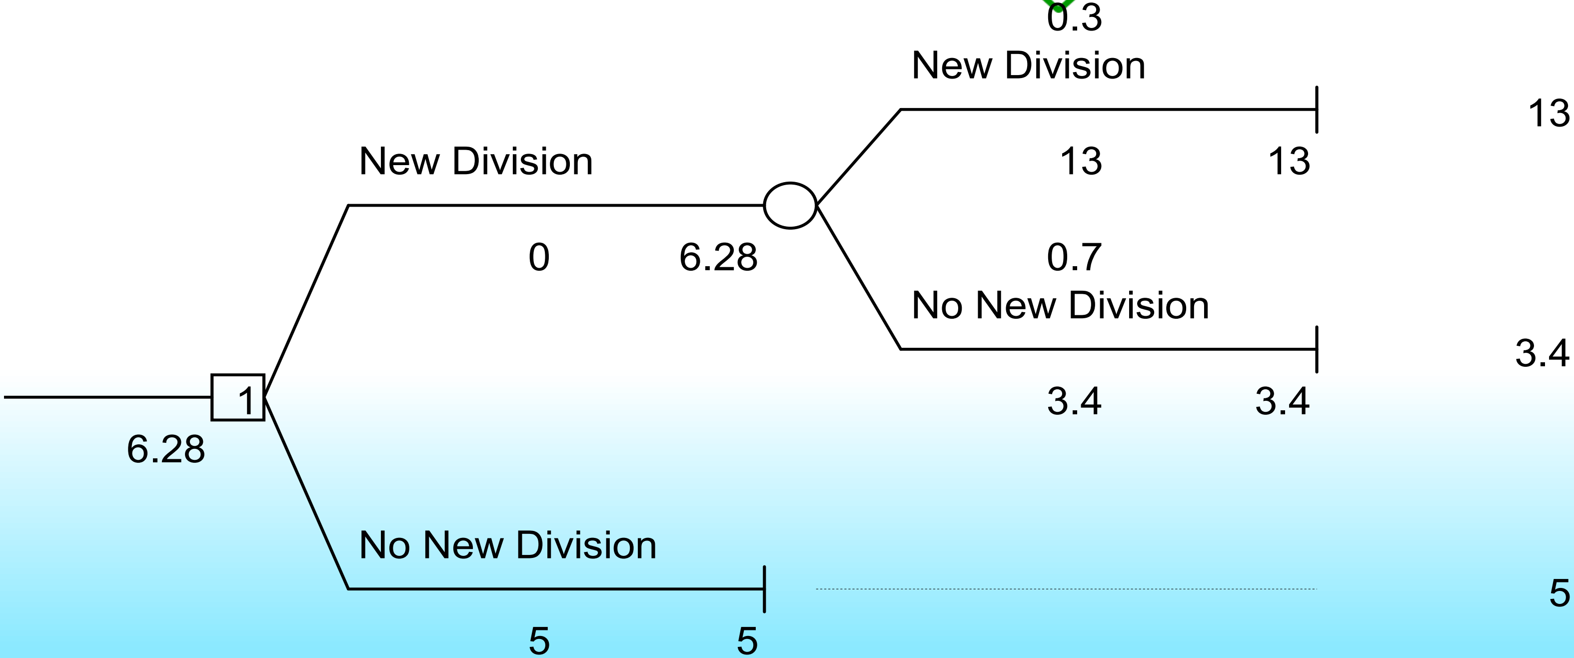

EVPI

- **Suppose** You were offered the possibility to come to know whether the **competitor** will enter the **market** for sure.
- Would You **pay** for that **info**? (i.e., do You think it **adds value** to the decision?)
- **Yes**.
- So, **how much** would **You pay**?

The Expected Value of Perfect Info

The EVPI is the difference between:

- *T1: Maximum expected payoff that can be generated* ***with free, perfect information***

–

- *Maximum expected payoff that can be generated* ***without the additional info***

Computing the EVPI

- First the term **T1:** *Value of the decision problem with perfect info*
- In terms of ID, if You have **perfect information,** it is as if You knew whether the event **has happened or not**

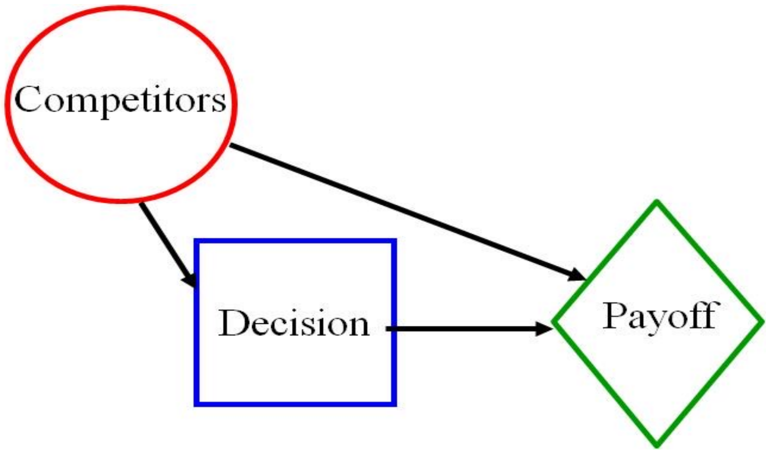

The corresponding DT

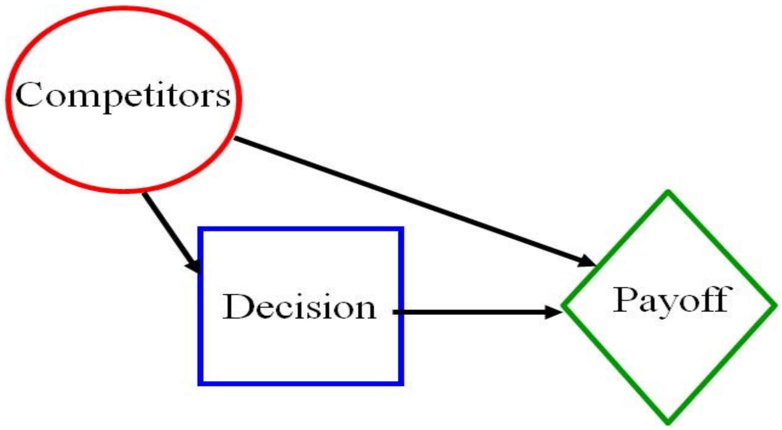

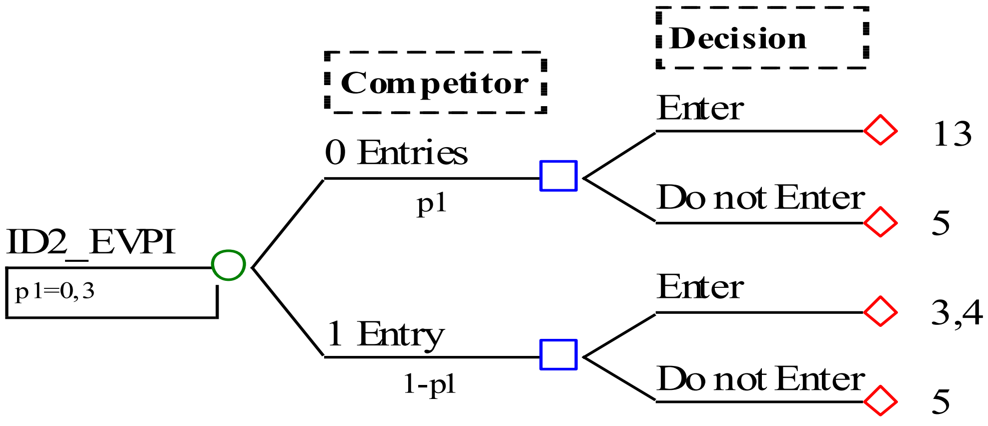

T1: Solve the EVPI DT

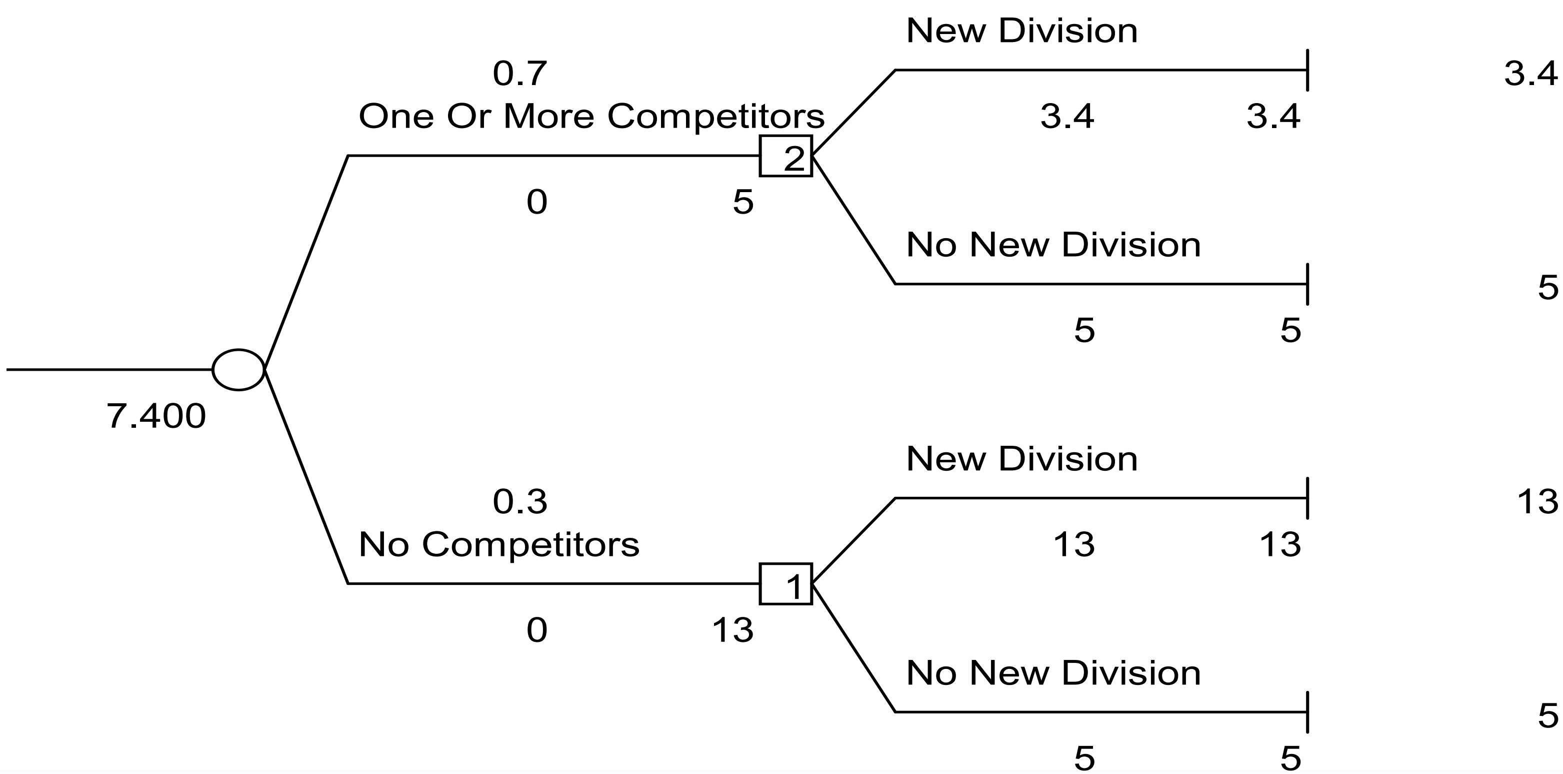

T1 and T2

- T1=7.4
- T2=6.28

EVPI=10.6-6.28=1.12

Expected Profit	6.28
Expected Profit with Perfect Info	7.40
EVPI	1.12

What Happens when info is not perfect?

- Expected Value of Sample Information

Bocconi

Università Commerciale Luigi Bocconi

Valuing New Information

- Should **BB hire Inside Games**, a market research company, who will **provide more information** about the likelihood of the **competitors** entering the **market**?
- Inside Games, offers the **following deal**.
- Inside Games can forecast whether it is **"likely"** or **"unlikely"** that new **competitors** will enter.
- ***If You want this information, You have to pay*: $0.3 million!**
- ***Quality*** of the information is an **important factor**.

Quality of Information

- BB's opinion regarding the **quality** of the information provided by Inside Games is captured through the specification of ***likelihood probabilities*** that are data of the problem:
  - the likelihood that, ***if*** competitors actually **enter** the market, then Inside Games supplies a **"Likely"** report.
  - the likelihood that ***if*** competitors actually do ***not*** **enter** the market, then Inside Games supplies an **"Unlikely"** report.

- ***Accuracy of the source!!!!***

ID with Imperfect Info

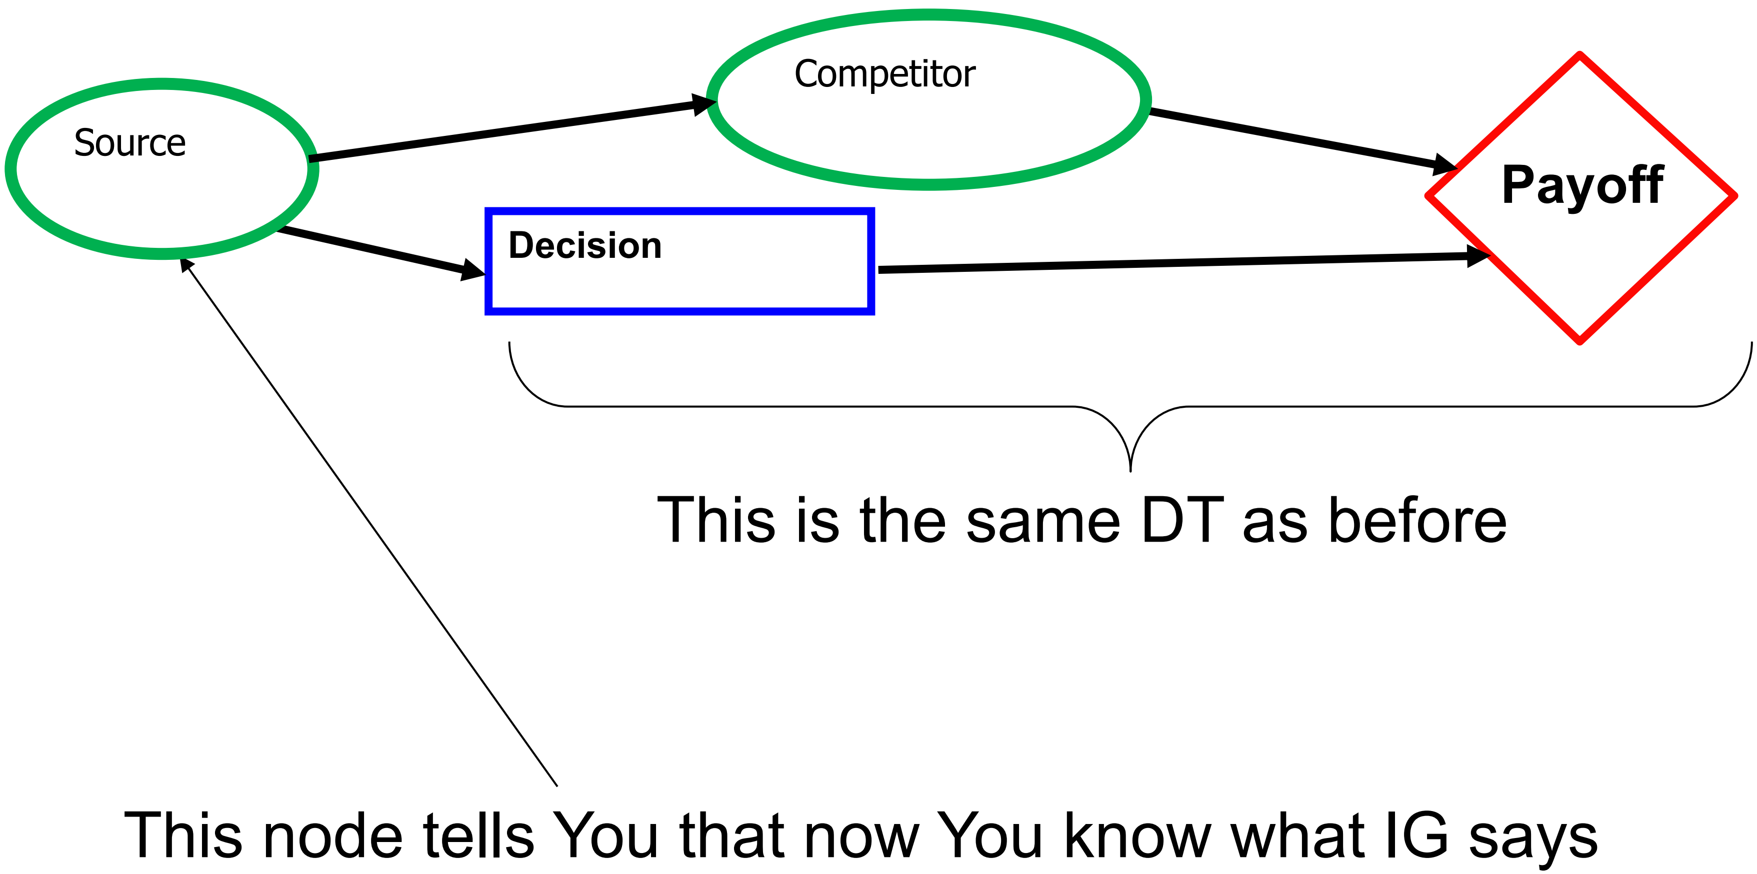

The EVSI Decision Tree

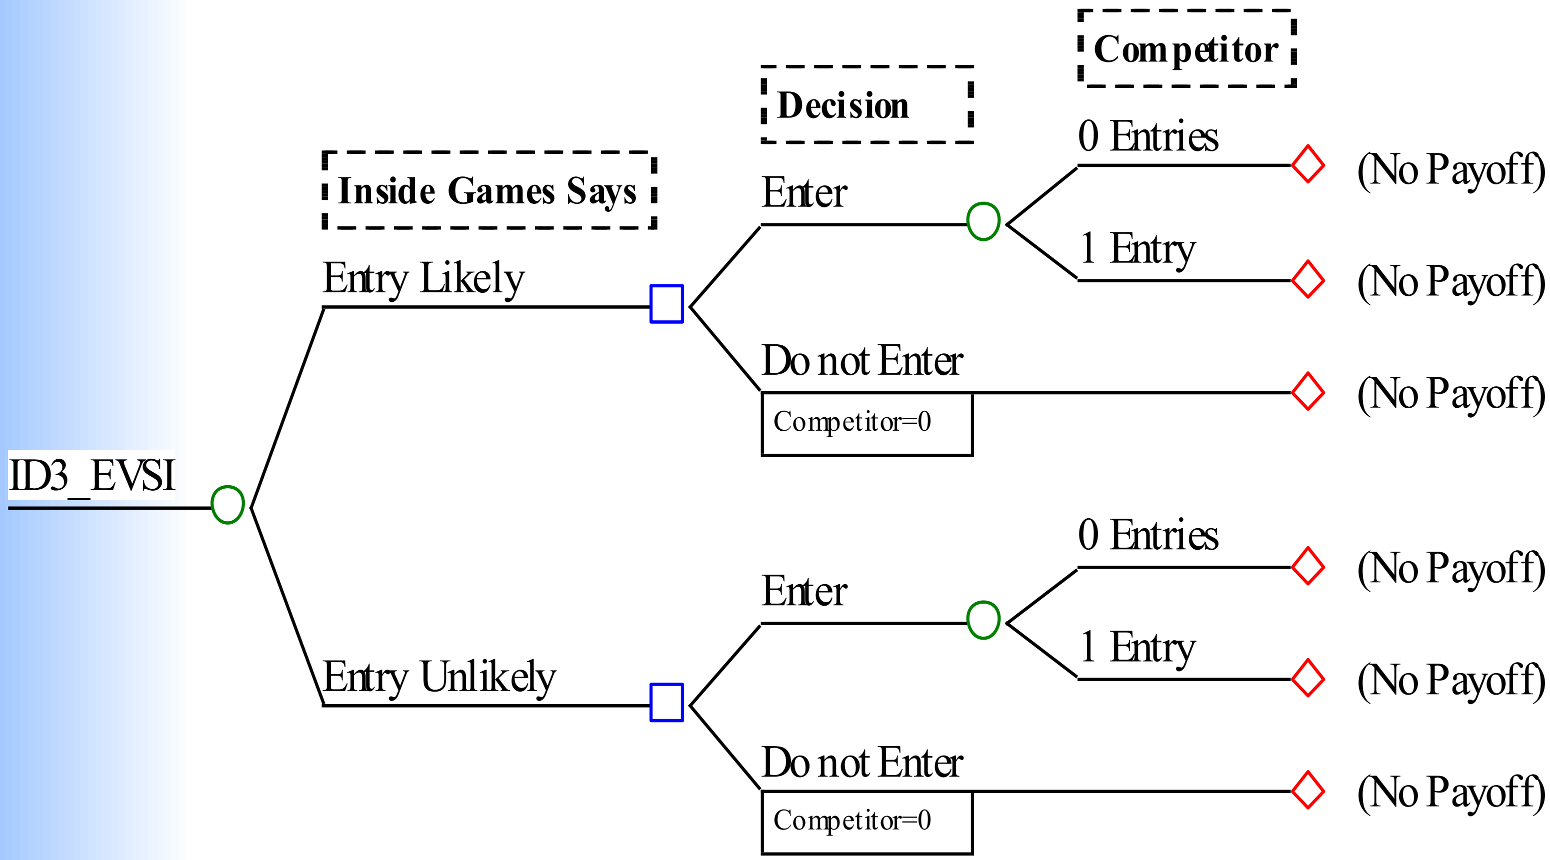

EVSI

- *The **Expected Value of Sample Information** (EVSI) is the difference between:*

- *The Maximum expected payoff that can be generated with **free, imperfect information***

- *The Maximum expected payoff that can be generated with **no information***

Calculation Procedure

- First **Find T1** and then **T2**
- However, here there are **some differences**:
- Note that in the EVPI case we did not have any **chance node** between the **decision** and the **payoff**
- Now, we have the **event node "source"** before the decision node and the **event node "competitors"** between the decision node and the terminal node.

EVENTS

Inside Games says likely (SL)

Inside Games says unlikely (SU)

Competitors enter the market (E)

Competitors do not enter market (D)

Unknowns

- ***We need to determine:***
- ***P(E|SL)***
- ***P(D|SL)***
- ***P(E|SU)***
- ***P(D|SU)***
- ***P(SL)***
- ***P(SU)***

Inserting in the Tree

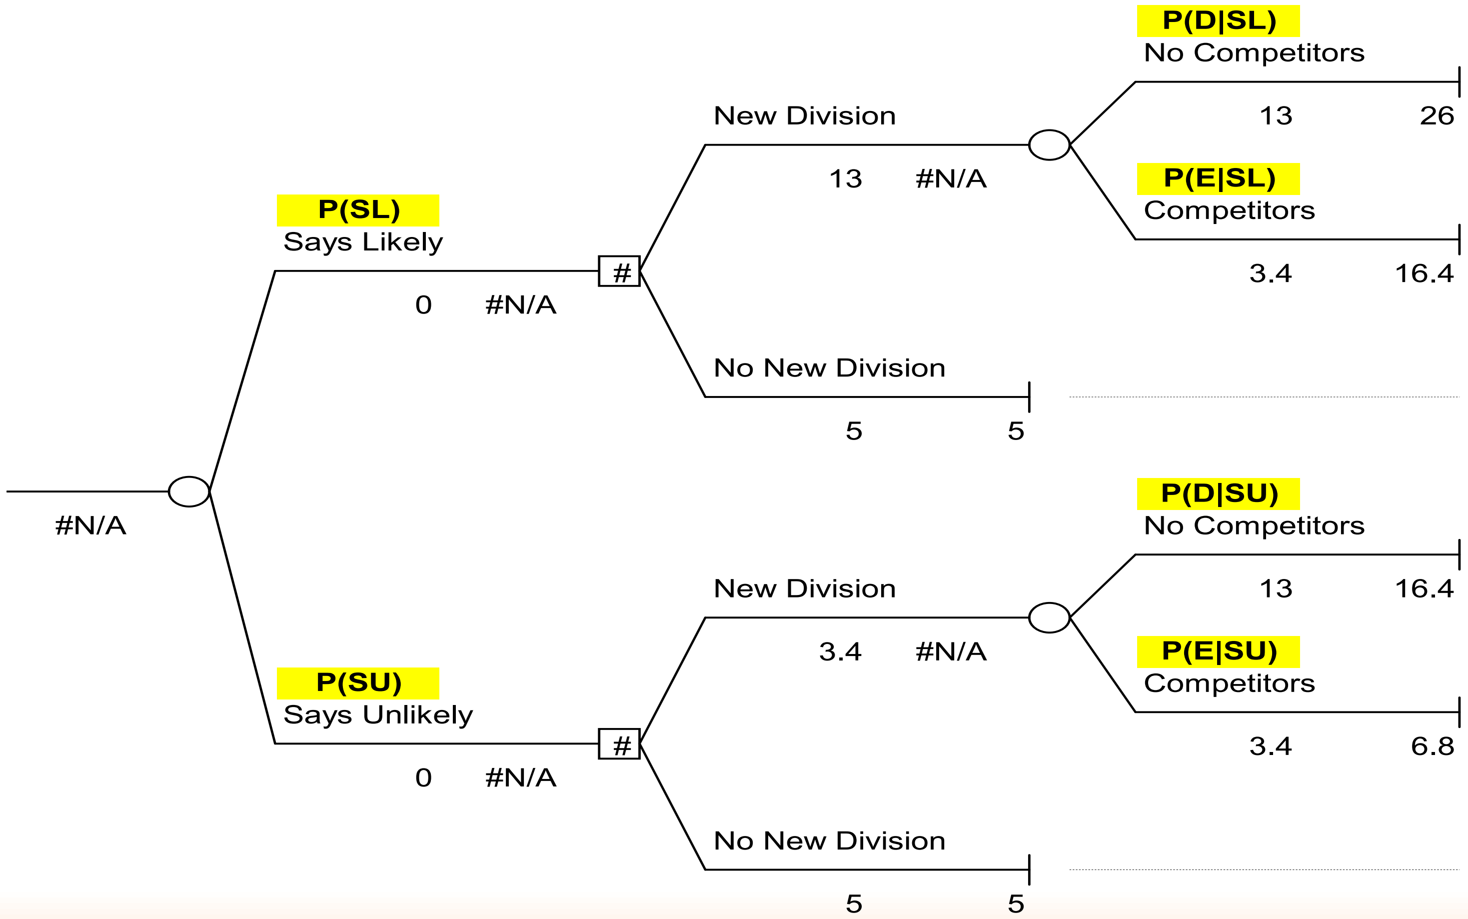

Bocconi

Università Commerciale Luigi Bocconi

DATA

- Recall that: Prior to **purchasing information** from Inside Games, **BB believes** that there is:
- a **30%** chance that **competitors enter** the market;
- a **70%** chance that there will be **no competitors**
- The **accuracy** of Inside Games is **decribed** by the following **Probabilities**
- $P(SL|E) = 0.75$
- $P(SU|D) = 0.75$.

A further step

- Now **suppose** Inside Games sends a **"Likely" report.**
- **Should we change** our view of the problem?
- What is **the probability** that **A-Starry** enters? (Higher or lower than 0.35?)
- We need to **determine *P(E/L)*,** i.e., the new probability that **A-Starry enters** the market ***if*** we have received a **"Likely"** report from Inside Games.
- Bayes' Rule (see Chapter 8).

Bocconi

Note

- In the tree, we need to **insert** ***P(E|SL)***
- However, we have data on ***P(SL|E)***
- ***How do we switch from one to the other?***

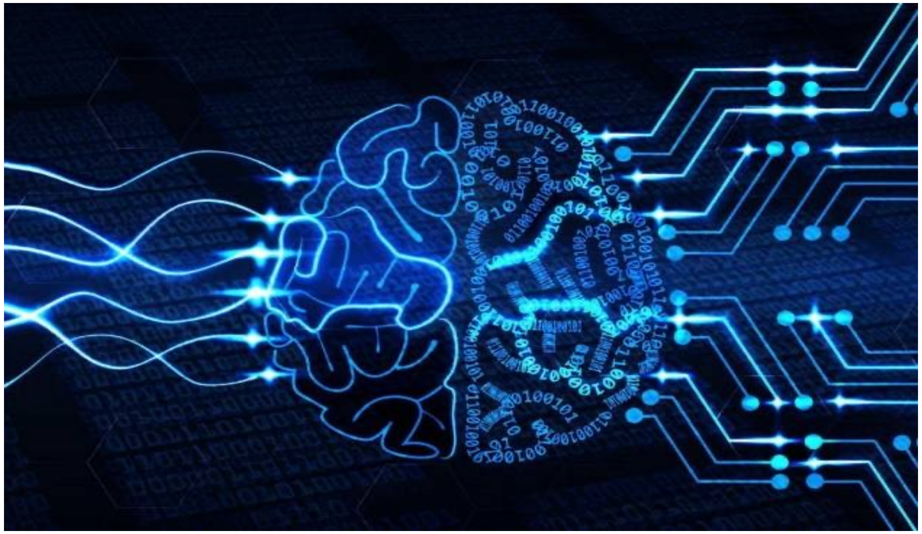

Bocconi

Bayes' Theorem

- **A e B** are two events. Event **A has happened**.
- **Thesis:** the probability of **B** given that A has happened is **updated** as follows

Prior Probability of B

Likelihood

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

Posterior Probability of B

Probability of A

The total probability rule

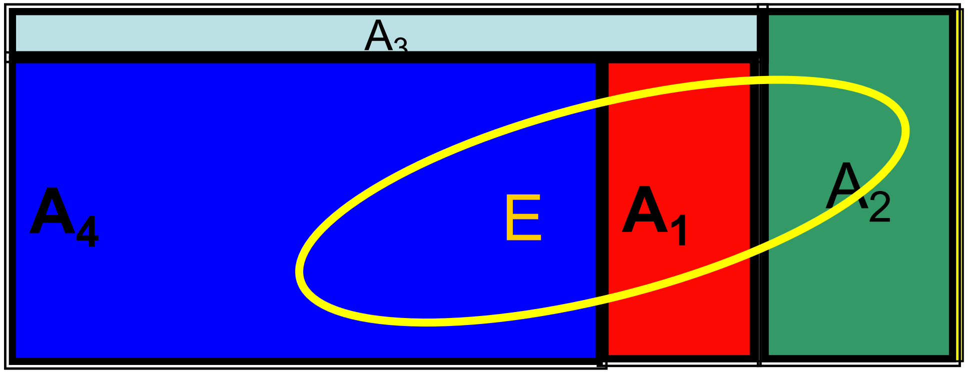

- N mutually exclusive events (A_1, A_2,…,A_N) and exhaustive:

$$1 = P(A_1) + P(A_2) + \ldots + P(A_N)$$

- Then :

$$P(E) = P(E \cap A_1) + P(E \cap A_2) + \ldots + P(E \cap A_N) =$$

$$= P(E \mid A_1)P(A_1) + P(E \mid A_2)P(A_2) + \ldots$$

Bayes Theorem

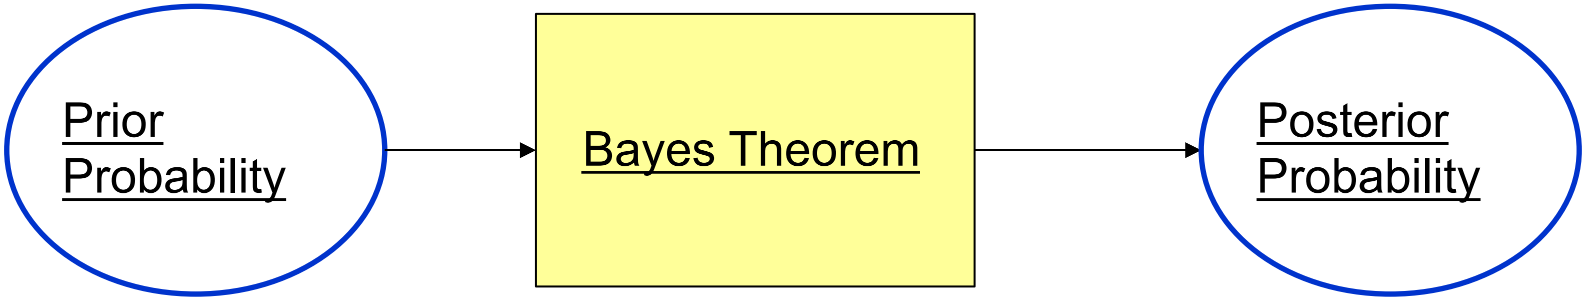

Prior Probability: Value of the Probability Before Hearing the Source

Posterio Probability: Value of the Probability after Hearing the Source

In Our Example

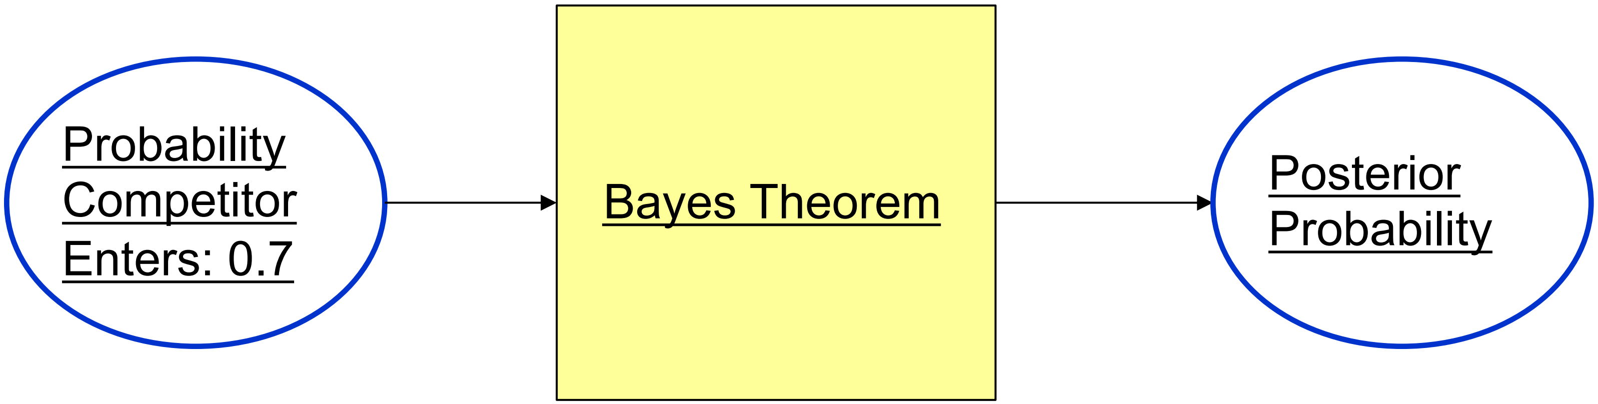

Posterior Probability I: P(E|+)

Probability Competitor Enters Given that the Source Says Enter

Posterior Probability II: P(E|-)

Probability Competitor Enters Given that the Source Says Does not Enter

Bocconi

Bayes File Calculates The Probabilities

Accuracy I

Prior Probability

Pr(E)	**Pr(+|E)**	**Pr(-|nonE)**
0.7	0.95	0.85
Pr(nonE)	**Pr(-|E)**	**Pr(+|nonE)**
0.3	0.05	0.
Pr(+)	**Pr(E|+)**	**Pr(E|-)**
0.71	0.93661972	0.12068966
Pr(-)	**Pr(nonE|+)**	**Pr(nonE|-)**
0.29	0.06338028	0.87931034

Accuracy II

Probability Says Comp. Enters

Posterior II

Posterior I

In Our Example

Probability Competitor Enters	Probability Says "Enters" | Competitor Enters	Probability Says "Does not Enter"| Competitor Does not Enter
P(_E)	Pr(+|E)	Pr(-|nonE)
0.7	0.95	0.85

EVSI for the Example

- Value with the imperfect information: 6.984
- Value without information: 6.28
- Expected Value of Sample Information:
- EVSI=6.474-6.28=0.704

Additional Slides with different Data

For Excercise

Bayes' Theorem

$$P(E \mid SL) = \frac{P(E) \cdot P(SL \mid E)}{P(SL)}$$

0,7

0,75

Unknown, but…

The total Probability Rule

- $P(SL) = P(SL \cap E) + P(SL \cap D)$
- $= P(SL|E) \times P(E) + P(SL|D) \times P(D)$
- $= 0.75 \times 0.7 + (1 - 0.75) \times 0.3$
- $= 0.6.$

- *In addition P(SU)=?*
- *Well, P(SU)=1-P(SL)=0.4*

Bayes' Theorem

$$P(E \mid SL) = \frac{P(E) \cdot P(SL \mid E)}{P(SL)}$$

0,7

0,75

Unknown, but…

$$P(E \mid SL) = 0.8750$$

Bayes' Theorem

$$P(E \mid SU) = \frac{P(E) \cdot P(SU \mid E)}{P(SU)}$$

0,7

0.4

0,25

$$P(E \mid SU) = 0.4375$$

Bocconi

The Other Missing Data

- *We need to determine:*
- *P(E|SL)=0. 8750*
- *P(D|SL)=1-P(E|SL)= 0.125*
- *P(E|SU)=0.4375*
- *P(D|SU)=1-P(D|SL)=0. 5625*
- *P(SL)=0.6*
- *P(SU)=0.4*

The Completed Tree

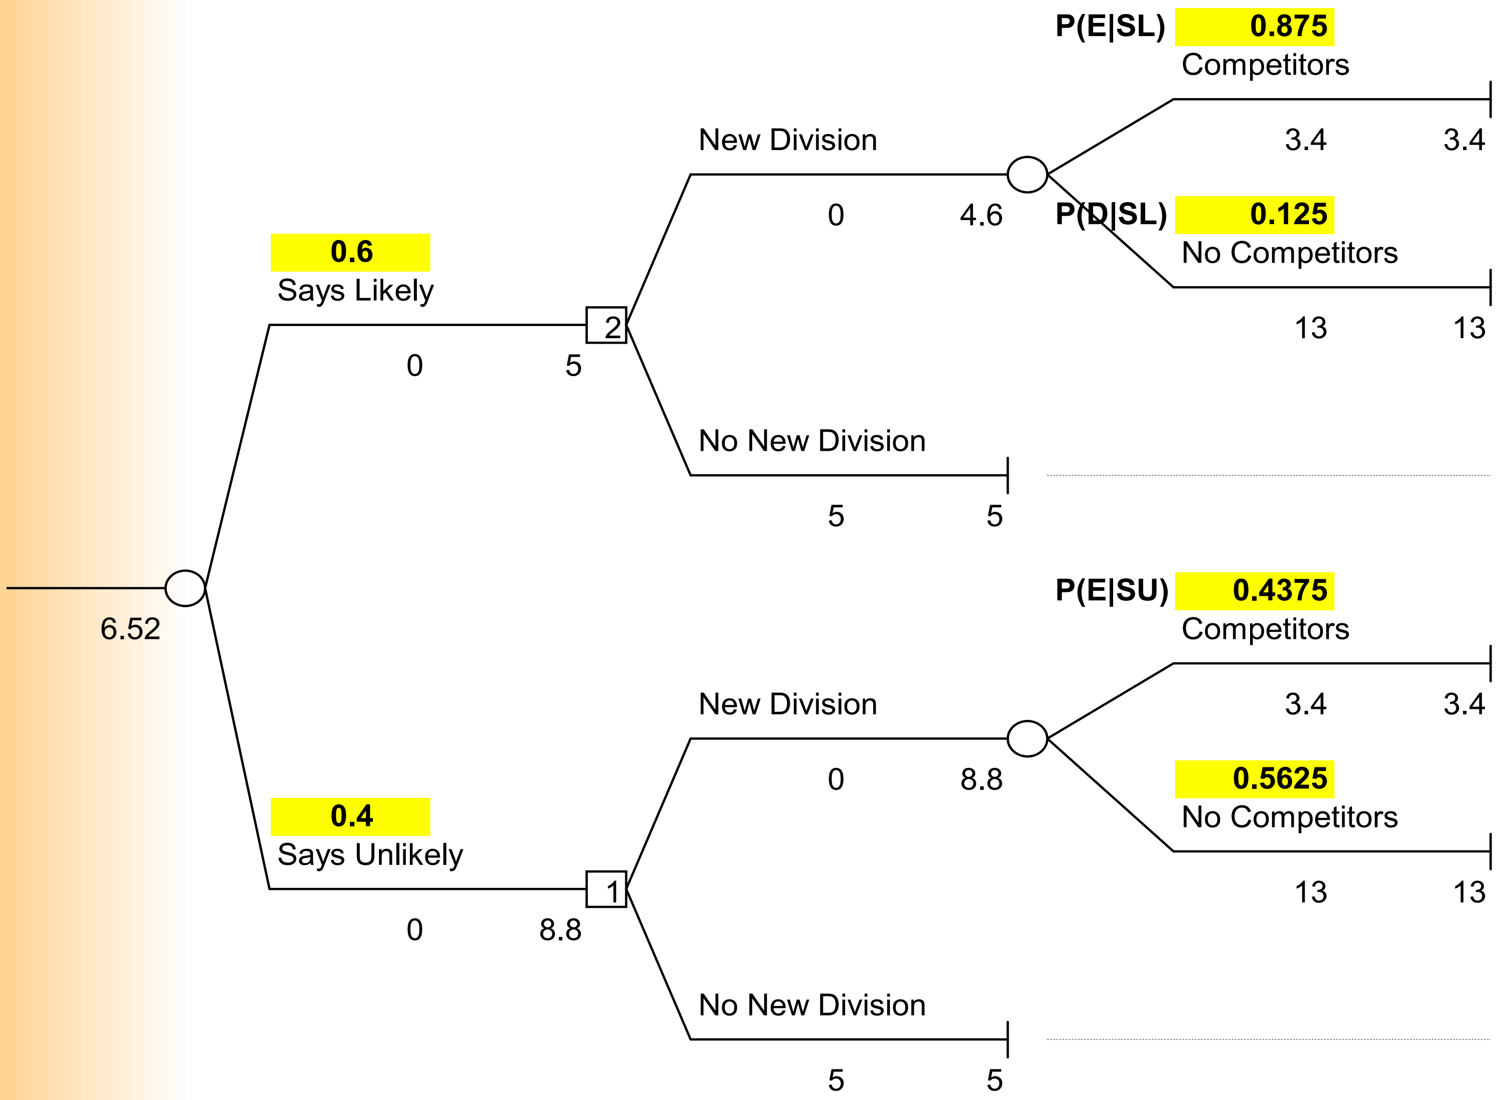

Final Result

Expected Profit	6.28
Expected Profit with Imperfect Info	6.52
EVSI	0.24

Bocconi

End of Additional Slides with different Data

For Excercise

In summary

Does the new info add value?

Utility and Preference

- **Maximizing** expected payoffs ignores issues associated with ***risk*.**
- Consider the **game of flipping a fair coin for $500**: Heads you win; tails you lose.
- **How much** must you be **paid** in order to play this game once?
- If this amount is **positive**, you are ***averse to risk***.
- We now study ***risk aversion*** and the impact that it has on **decision–making**.
- **The difference**

$$R = E[X] - z$$

is called ***Risk Premium***

Risk

The **Utility function** incorporates **subjective attitudes** towards **risk**

The objective goes from

maximize expected profit

to

maximize expected utility

Utility Function

It is built interpolating from:

1. **Non-satiation: increasing**
2. **E[Lottery]= *z***

Attitudes towards risk

Attitude	Utility Function	Risk Premium
Neutral	Linear	0
Averse	Concave	R>0
Prone	Convex	R<0

Arrow-Pratt Risk Aversion Function:

$$r(x)=-u''(x)/u'(x)$$

Exponential Utility Function and Certainty Equivalent

- Let **x** be a **monetary amount**. Then, the exponential utility of x is defined as

$$u(x) = 1 - e^{-\lambda x}$$

- The **larger λ** the **larger the DMr risk aversion**
- Constant risk aversion: ***r(x)=λ***
- Let **E[U(x)]** be the **expected utility** of the **lottery**. Then, the **certainty equivalent** of the lottery is

$$z = -\frac{1}{\lambda}\ln[1 - E[U(x)]]$$

A More General form

More in general, we can write:

$$u(x) = A - Be^{-x/RT}$$

with A and B and RT positive. Setting:

$$A = 1$$

$$B = 1$$

$$RT = 1/\lambda$$

we find back the formula in the previous slide.

The **value RT** is called **"risk tolerance",** and it is the **inverse** of the **risk aversion constant**

Node Valuation

Node Type	Value Assignment Rule
Event Node	Expected utility of the branches emanating from the node. One can also compute the certainty equivalent
Decision Node	<u>Max</u> of the expected utilities of the emanating branches.

Bocconi