ECO2003F

Katherine Eyal

Lectures 17 - 32

Chapters 5/6, 7, 8, 12.

Chapter 5/6 - lectures 17, 18, 19, 20, 21, 22

My expectations

Read the textbook section **before** the lecture

Come to class ready to listen and engage with the material

Stop me and ask me to explain something if you don't understand

Let me know if something needs further revision

Interact with me and each other respectfully

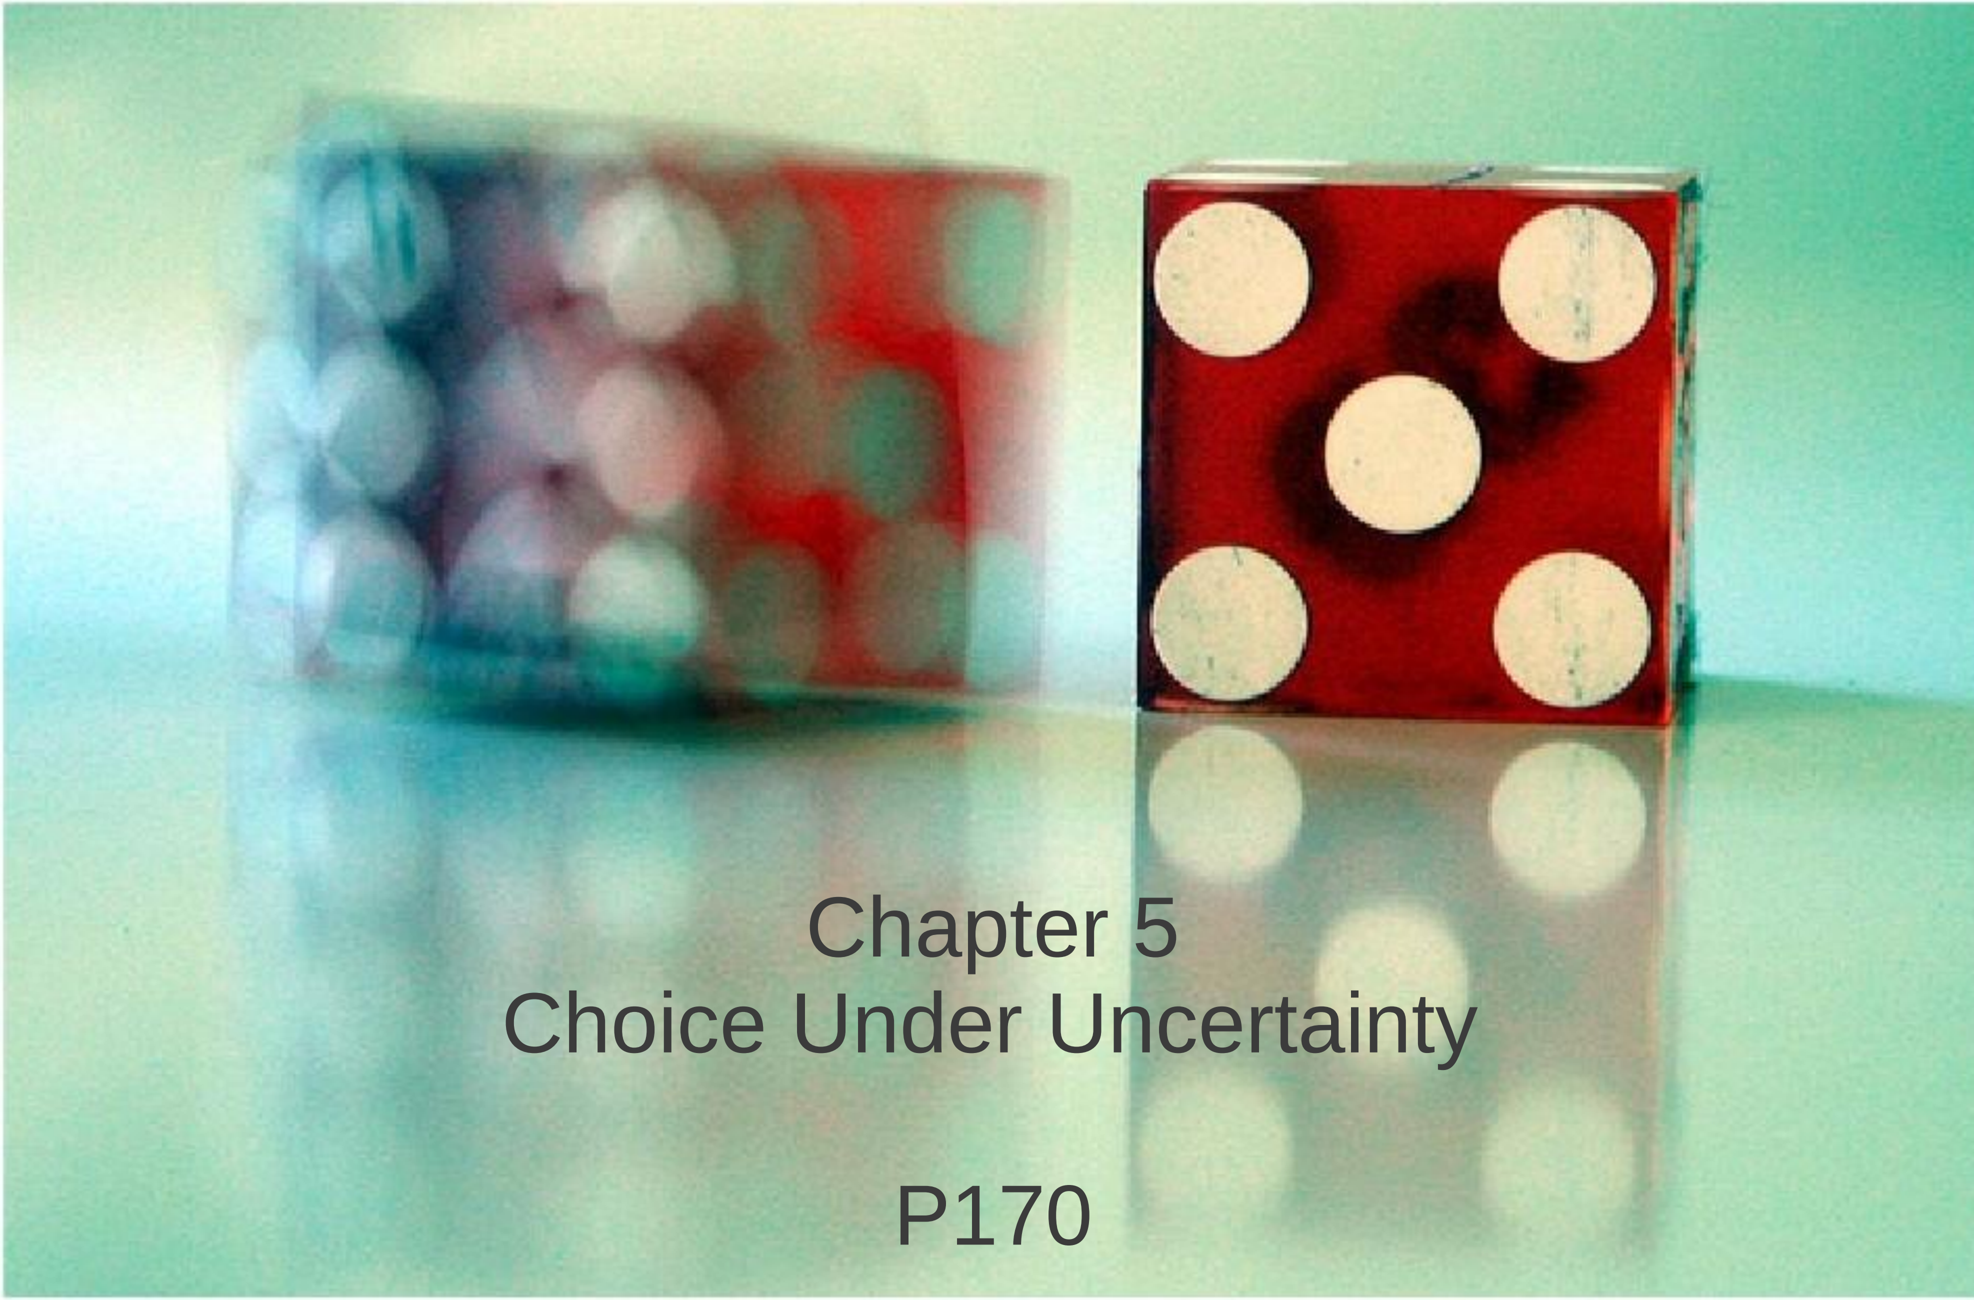

Chapter 5
Choice Under Uncertainty

P170

Do we have **complete knowledge** of all the alternatives when making a choice?

Which new phone should we buy?

Do you arrange an outdoor picnic in May?

?

All decisions involve risks.

Decision making involves trying to understand these risks.

Probability and Expected Value

A decision under uncertainty is a **gamble**

E.g. Throwing a **fair** coin:

1: if heads, win R1000, if tails, lose R5
2: if heads, win R2000, if tails lose R1000
3: If heads, win R200 000, if tails, lose R100 000

Which gamble do we take? How to decide?

Expected Value: weighted average of outcomes

Expected Value of a gamble Y with:

10 **outcomes**:

$Y_1, Y_2, \ldots, Y_{10}$

the **probabilities** of each occurring:

$a_1, a_2, \ldots, a_{10}$

Then the expected value of the gamble Y is:

$$\mathbf{E(Y) = a_1 Y_1 + a_2 Y_2 + \ldots + a_{10} Y_{10}}$$

Definitions:

a fair coin:
probability of heads is 50%, and tails 50%

If we throw a fair coin many times, we expect to get heads 50% of the time

Fair bet:
Expected Value = 0

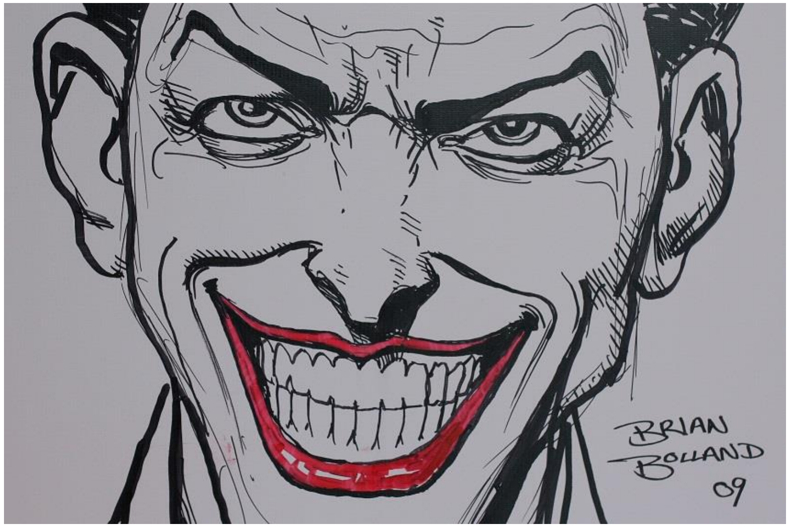

E.g. Throwing a **fair coin:**

Probability of heads is 50%, and tails 50%

1: if heads, win R1000, if tails, lose R5
2: if heads, win R2000, if tails lose R1000
3: If heads, win R200 000, if tails, lose R100 000

E(1) = 0.5*(1000) + 0.5*(-5) = 497.5
E(2) = 0.5*(2000) + 0.5*(-1000) = 500
E(3) = 0.5*(200 000) + 0.5*(-100 000) = 50 000

Do we agree to 1? And to 2 and 3?
3 looks so attractive - why don't we take it?

What makes a gamble attractive?

Is a **positive expected value** enough to make you take a gamble?

Why/why not?

What else matters in the decision?

?

Von Neumann-Morgenstern Expected Utility Model

We choose the alternative
with the highest **expected utility,**
not the highest **expected value**

http://www.atomicarchive.com/Bios/vonNeumannPhoto.shtml

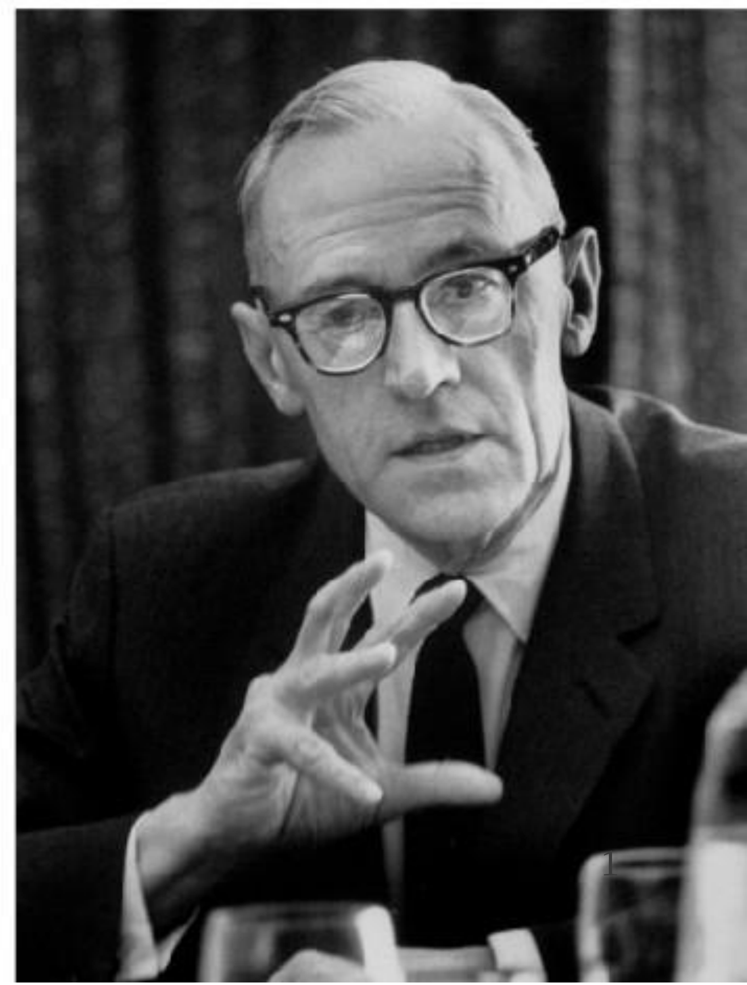

http://www.biografiasyvidas.com/biografia/m/fotos/morgenstern.jpg

We assume a **utility function U**
U assigns a number to each outcome

We start with **initial wealth M_0**

EU = expected value of utility over all possible outcomes:

Expected $U_1 = 0.5*U(M_0 + 1000) + 0.5*U(M_0 - 5)$

Can you work out the expected utility of 2 and 3?

E.g. 5.1 on page 173

$$U(M) = \sqrt{M}$$

$$EU_1 = 0.5\sqrt{101000} + 0.5\sqrt{99995}$$

$$EU_2 = 0.5\sqrt{102000} + 0.5\sqrt{99000}$$

$$EU_1 = 0.5\sqrt{300000} + 0.5\sqrt{0}$$

Expected **values** of a set of outcomes **may not have the same ranking** as expected **utilities**.

Why not ?

E_1 = 497.5 , E_2 = 500, E_3 = 50 000

Initial wealth 100000, & utility fn $U(M) = \sqrt{M}$

EU_1 = 317.012, EU_2 = 317.009, EU_3 = 273.861

Which option is most attractive?
Is this more logical?

Why are the answers so different when we compare expected value and expected utility?

Which makes more sense?

Utility functions are often **non linear**.

What shapes are common for utility functions?

?

Figure 5.2: A Concave Utility Function

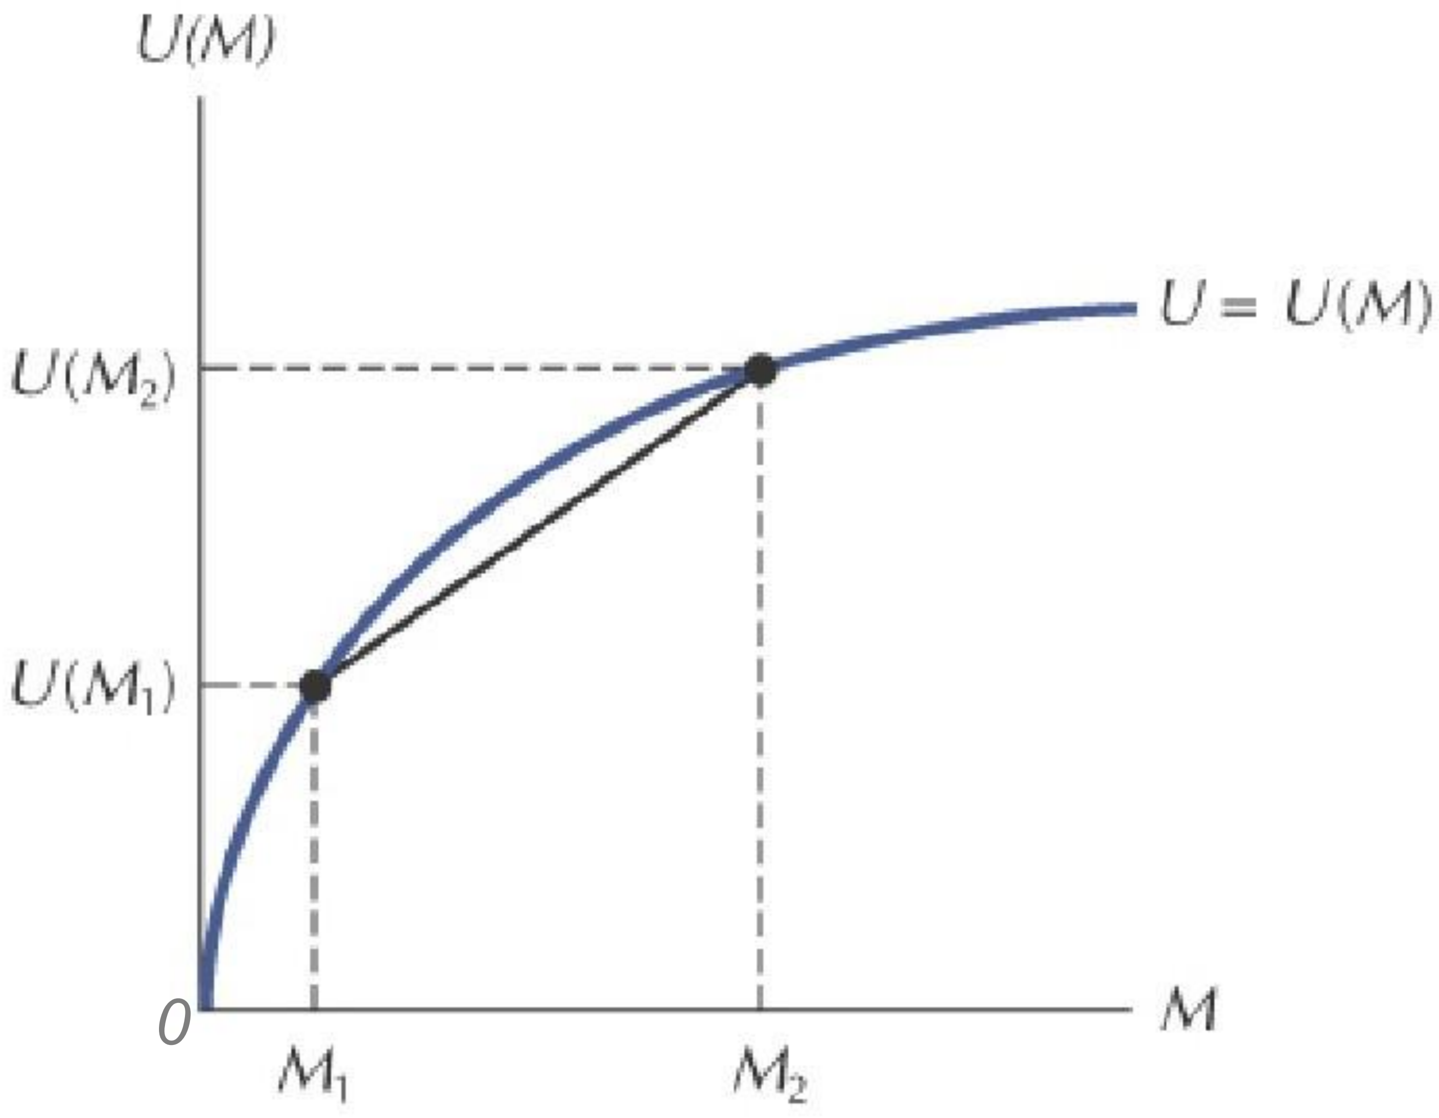

Concave Utility Functions

NB: A function is **concave** if for any values M_1 and M_2, the function lies **above** the chord joining the points [M_1, $U(M_1)$], [M_2, $U(M_2)$]

Exhibit **diminishing marginal utility** of wealth

The more wealth a consumer has, the lower the marginal utility from a one unit increase in wealth

Describes a **risk averse** consumer

Types of Risk Preference

Type	Risk Averse	Risk Neutral	Risk loving
Shape U function	Concave :(	Linear /	Convex :)
Marginal Utility	Diminishing	Constant	Increasing

When do we accept gamble 1?

?

If we don't accept it,

$EU = U(M_0)$

If we do accept it,

$EU = EU_1$

In general, we accept a gamble G if $EU_G > U(M_0)$

A **fair gamble** has expected **value** = 0

Fair Gamble G:

throw a coin, gain R30 if heads, lose R30 if tails

Expected Value = 0.5(R30) + 0.5(-R30) = 0

Bob has initial wealth = 40, utility fn U = U(M):

$$EU_G = 0.5U(40 + 30) + 0.5U(40 - 30)$$
$$= 0.5U(70) + 0.5U(10)$$

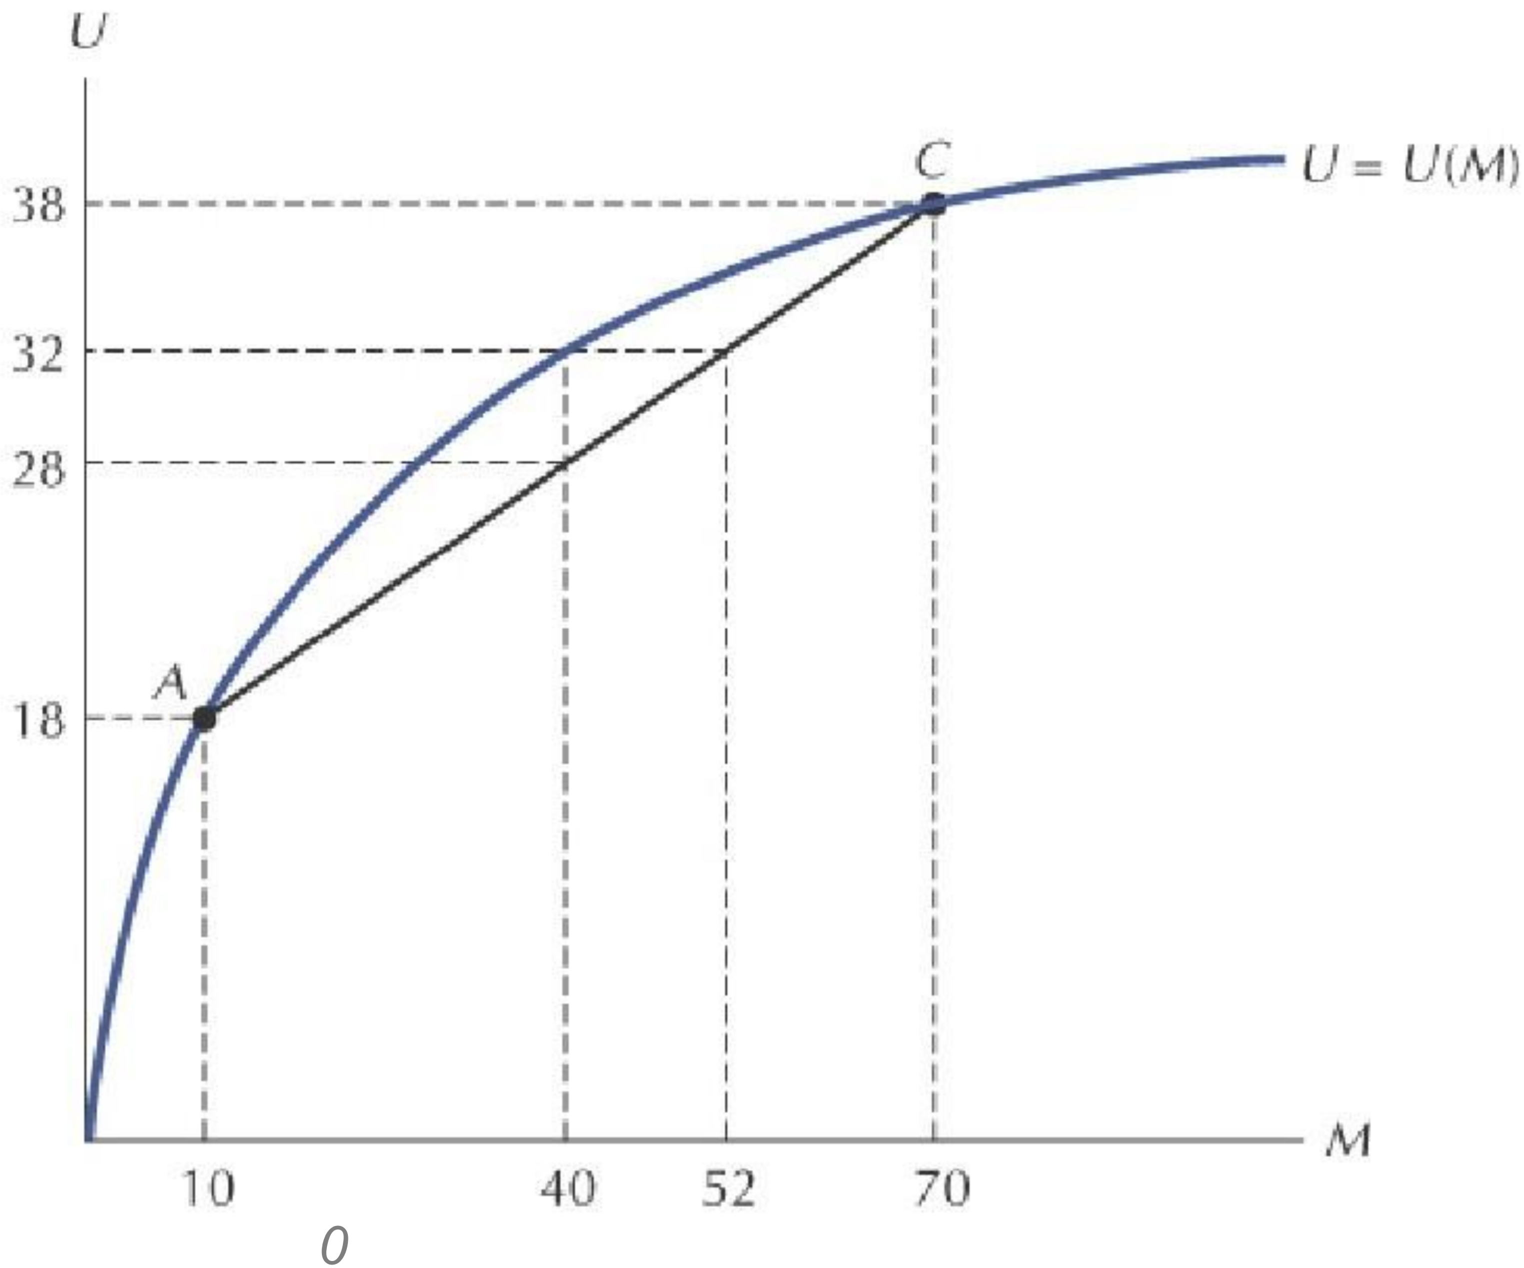

A Risk-Averse Person Will Always Refuse a Fair Gamble

How to interpret EU_G

Construct chord AC
EU_G lies somewhere on chord AC
EU_G lies halfway if probability of winning = 0.5

If we **refuse** gamble G, then EU = U(40) = 32
If we **accept** gamble G,
EU = 0.5U(70) + 0.5U(10)
EU = 0.5(38) + 0.5(18) = 28

So U(Refusing) = 32 > U(Accepting) = 28
We refuse to take the fair gamble.

You must be able to:

1. Find the **expected value** of a gamble
2. Find the **expected utility** of a gamble
3. Know the expected utility of **refusing a gamble**
4. Know whether we accept or refuse a gamble, by comparing the results of 1 and 2.
5. Know the maximum expected value of a gamble that the individual would refuse, given the shape of the utility function
6. Know if the person is **risk averse/seeking/neutral**

N.B.!

Page 176

Example 5.2

Exercise 5.1

Answers Page 196

Rule: if the probability of winning is p, and of losing is 1-p, then the expected utility lies a fraction of 1-p to the left of C (the winning endpoint).

A Risk-Seeking Person has Convex U fn
$EU_G > U(M_0)$

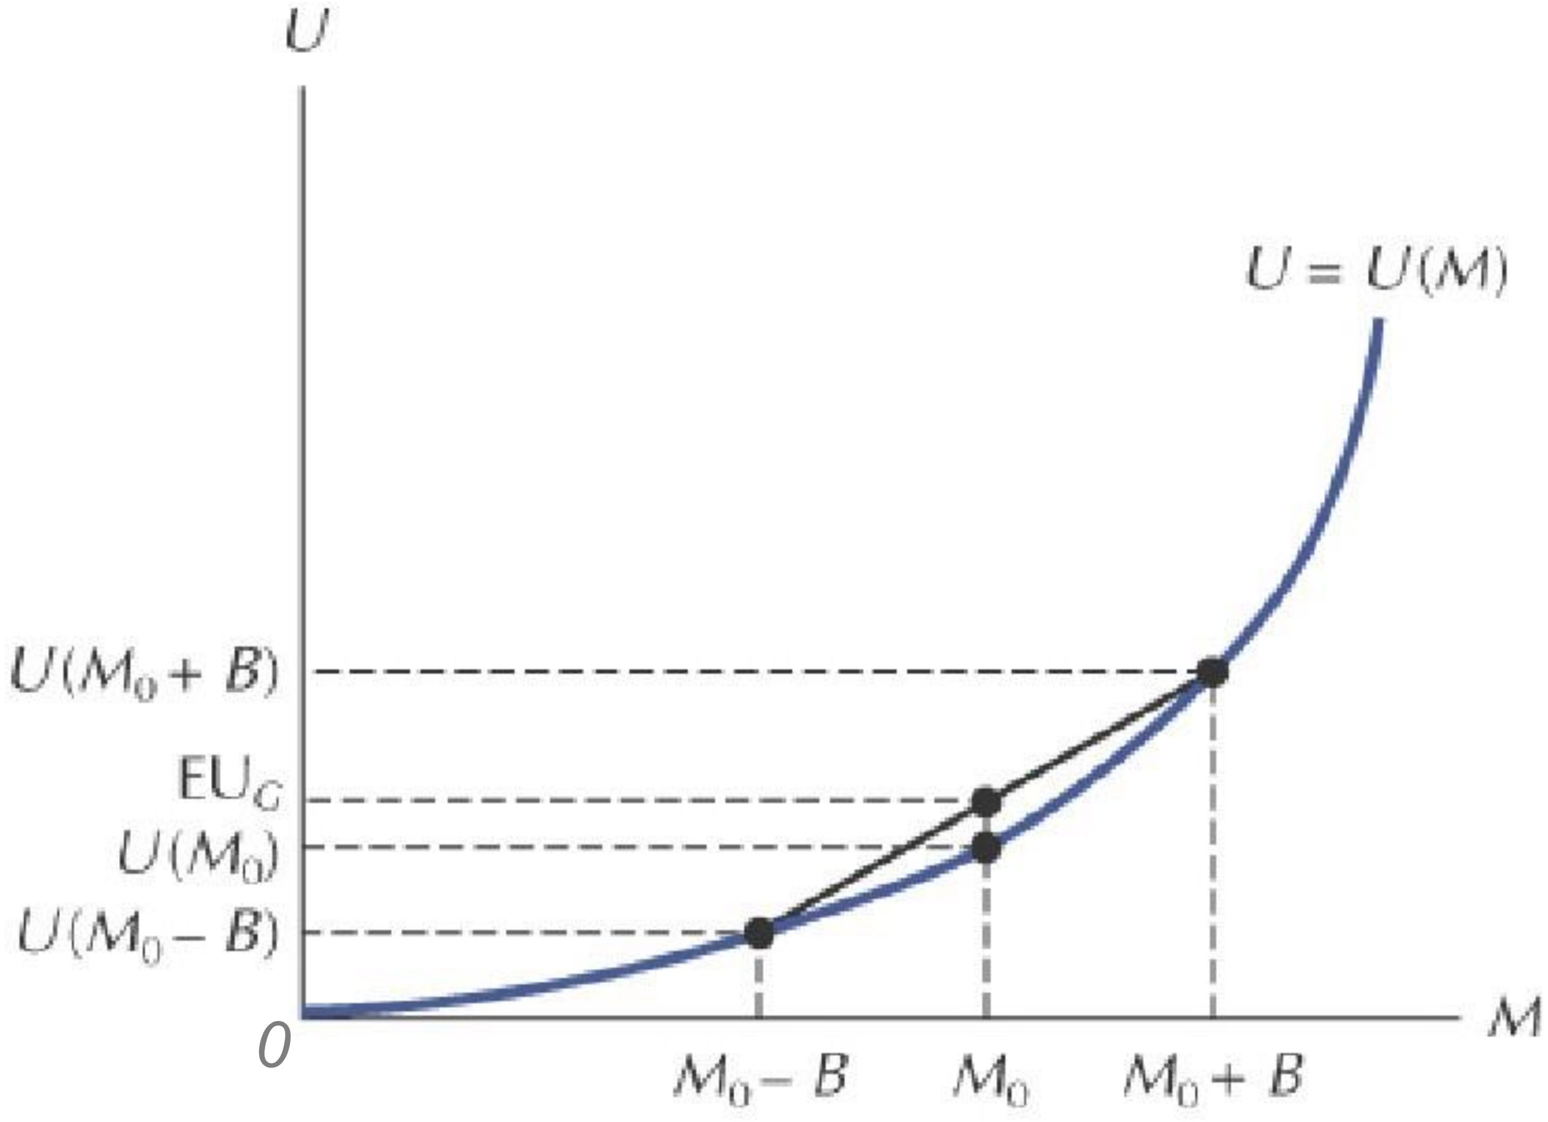

A Risk-Neutral Person has Linear U fn
$EU_G = U(M_0)$

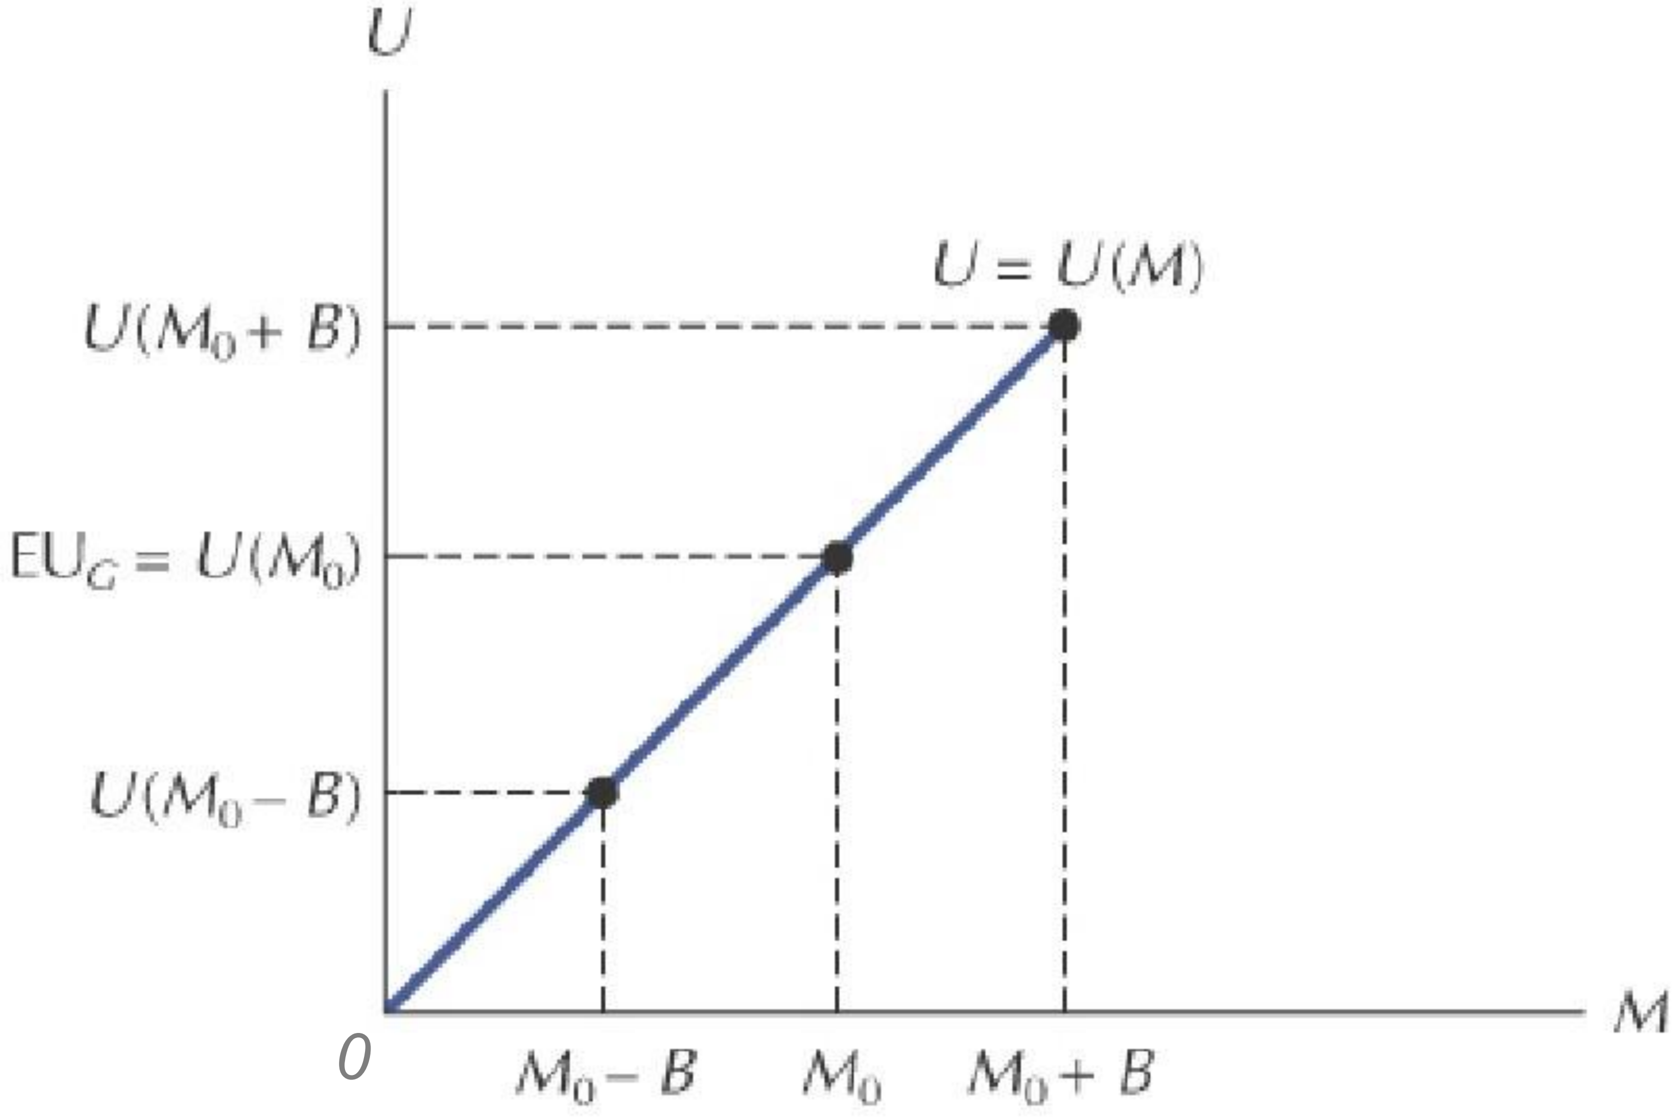

P176, Exercise 5.2, 5.3

Person with initial wealth 1000 can win 200 with probability 1/2, and lose 200 with probability 1/2. For utility functions $U(M) = M$, and $U(M) = M_2$, decide if she will take the gamble or not.

A gamble that has 3 **outcomes** B_1, B_2 and B_3,
Which occur with **probabilities** p_1, p_2 and p_3

has

$$EV = p_1B_1 + p_2B_2 + p_3B_3$$

$$EU = p_1U(B_1) + p_2U(B_2) + (1-p_1-p_2)U(B_3)$$

Why?

?

Are you going to be a **movie star**?!

Boingotlo has U fn U = 1-1/M

If she becomes a CA, M = 5 with p = 1

If she tries to be a star,
M = 400 with p = .01,
And if she fails
M = 2 with p = .99

Barry Ronge knows if she will succeed
How much would she pay him for this info?

http://www.flickr.com/photos/jblover130193/6790629812/sizes/l/in/photostream/

Without any information:

<u>If she becomes a **CA**:</u>

Certain U = 1-1/5 = 0.8

<u>If she tries to be a **star**:</u>

EU = 0.01U(400) + 0.99U(2)   (error on P179)

= 0.01(1-1/400) + 0.99(1-1/2)

= 0.505

Without the info,
she becomes a CA (Why?)

She can pay P to **get info**. She will find out:

Either: she is the next Charlize (p = 0.01)
She earns 400 (but has paid P)

Or: she will make a great CA (p = 0.99)
She earns 5 (but has paid P)

<u>Her expected utility is:</u>

$EU_1 = 0.01U(400 - P) + 0.99U(5 - P)$

What is the largest amount Barry can charge her?

Let's make her **indifferent** between finding out or not finding out, and see what P is.

if she doesn't find out

$U = 1-1/5 = 0.8$

if she finds out

$EU_1 = 0.01U(400 - P) + 0.99U(5 - P)$

$= 0.01[1 - 1/(400 - P)] + 0.99[1-1/(5 - P)]$

Equate these two, and work out $P = 0.0494$

What if he charges $P = 0.05$?

ttp://www.google.co.za/imgres?imgurl=http://www.capetowndailyphoto.com/uploaded_images/jammie_shuttle_uct_IMG_3205-786562.pg&imgrefurl=http://www.capetowndailyphoto.c

Example 5.5 on P180

Rashaad is a poor student who was run over by a Jammie shuttle. Should he sue UCT?

Things to consider:

The cost of a lawyer to sue UCT, the expected payout from the lawsuit, whether Rashaad is risk averse or not, if the suit has positive expected value etc.

Please go through this in detail on your own

Example 5.3, 177

Nice example but not enough detail: you may omit

Rashaad has U = 1-1/M, M_0 = 7

If he sues, wins 5 (p = 0.5) OR loses with p = 0.5,

He **has to pay the lawyer 2**

EV of suing = 0.5(5) + 0.5(0) = 2.5

Factor in cost of 2, then EV = 0.5

Does he sue?

EU = 0.5U(7 + 5 - 2) + 0.5U(7 - 2)

= 0.5(1-1/10) + 0.5(1-1/5) = 0.85

What is U if he doesn't sue?

U(7) = 1-1/7 = 0.86.

Does he sue?

EU(filing suit) > U (not filing suit)

EV> 0, but he is risk averse, and won't sue

How can we make him agree to sue?

Luckily, the lawyer is **risk neutral**.

He can assume the risk.

He charges **2 prices**:

F1 if Rashaad **wins**, and F2 if he **loses**

Rashaad would like to be **indifferent** between winning and losing

i.e. have same end expected wealth

Equate

$$7 + 5 - F1 = 7 - F2$$

Will the lawyer do this? Yes, if his costs are met

$$0.5F1 + 0.5F2 = 2$$

$$F1 = 4.5, F2 = -0.5$$

What is Rashaad's end wealth now?

Is this good?

Are the incentives aligned for the lawyer to do a good job?

We just showed that a risk neutral party can induce a risk averse person to engage in a transaction, by acting like an insurance company for them.

In risky transactions, a **risk neutral seller** can make the transaction more attractive for a **risk averse buyer** by acting as an insurance company.

Certainty Equivalent: The sum of money for which an individual would be **indifferent** between receiving that sum and taking the gamble

What is the CE in the following graph?

What was the CE for Rashaad?

?

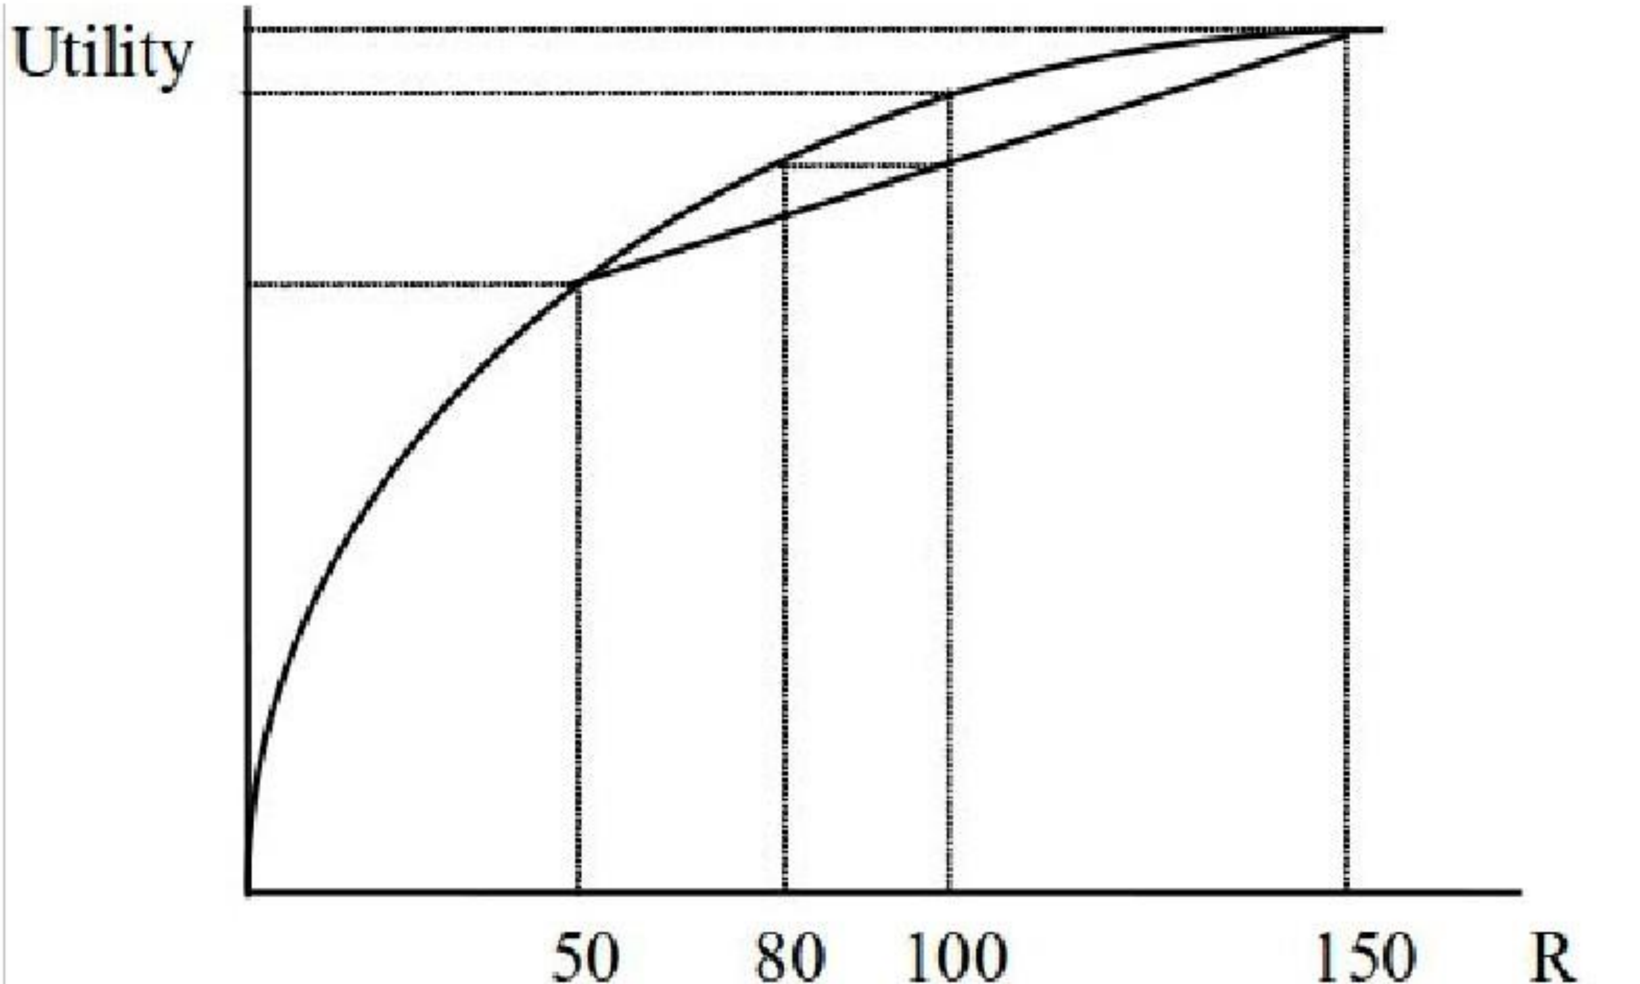

Figure 5-1

Starting income = R100

Bet:
lose R50 (p = 0.5), gain R50 (p = 0.5)

If **risk averse**, I don't like to gamble

In fact, I may pay you to not force me to gamble

I'll be happy to rather have R80, than gamble

So I'll pay R20 to avoid the gamble

R80 is the **certainty equivalent** of the gamble

0

I'm going to pay the insurance company R20, to keep the status quo

If I **lose**, the insurance company pays out R50 to get me back to R100

If I **win**, the insurance company pockets R50

Either way:
they get my R20, and I have R80 for sure

So they either lose R30, or they gain R70
Their **EV** = 0.5*(70) + 0.5(-30) = R20

0

Will you go into business as that insurance company?

If you had 1 client? 10000 clients?

We will return to this.

Insuring Against Bad Outcomes

If **risks are independent** for different consumers, it may be advantageous to practice **risk pooling**

3: If heads, win R200 000, if tails, lose R100 000

E_3 = R50 000, but no-one would agree to take it.

What if 10000 people agreed to each take the gamble, and share the proceeds evenly?

This works thanks to the **law of large numbers**

Law of **large** numbers

If an event happens independently with **probability p** in each of **N instances**, the proportion of cases in which the event occurs **approaches p,** as **N grows larger**.

E.g.

With a fair coin, as we throw the coin more and more times, the proportion of heads gets closer and closer to 50%.

H,H, T, H, T, T, T, H, H, T, H, T...

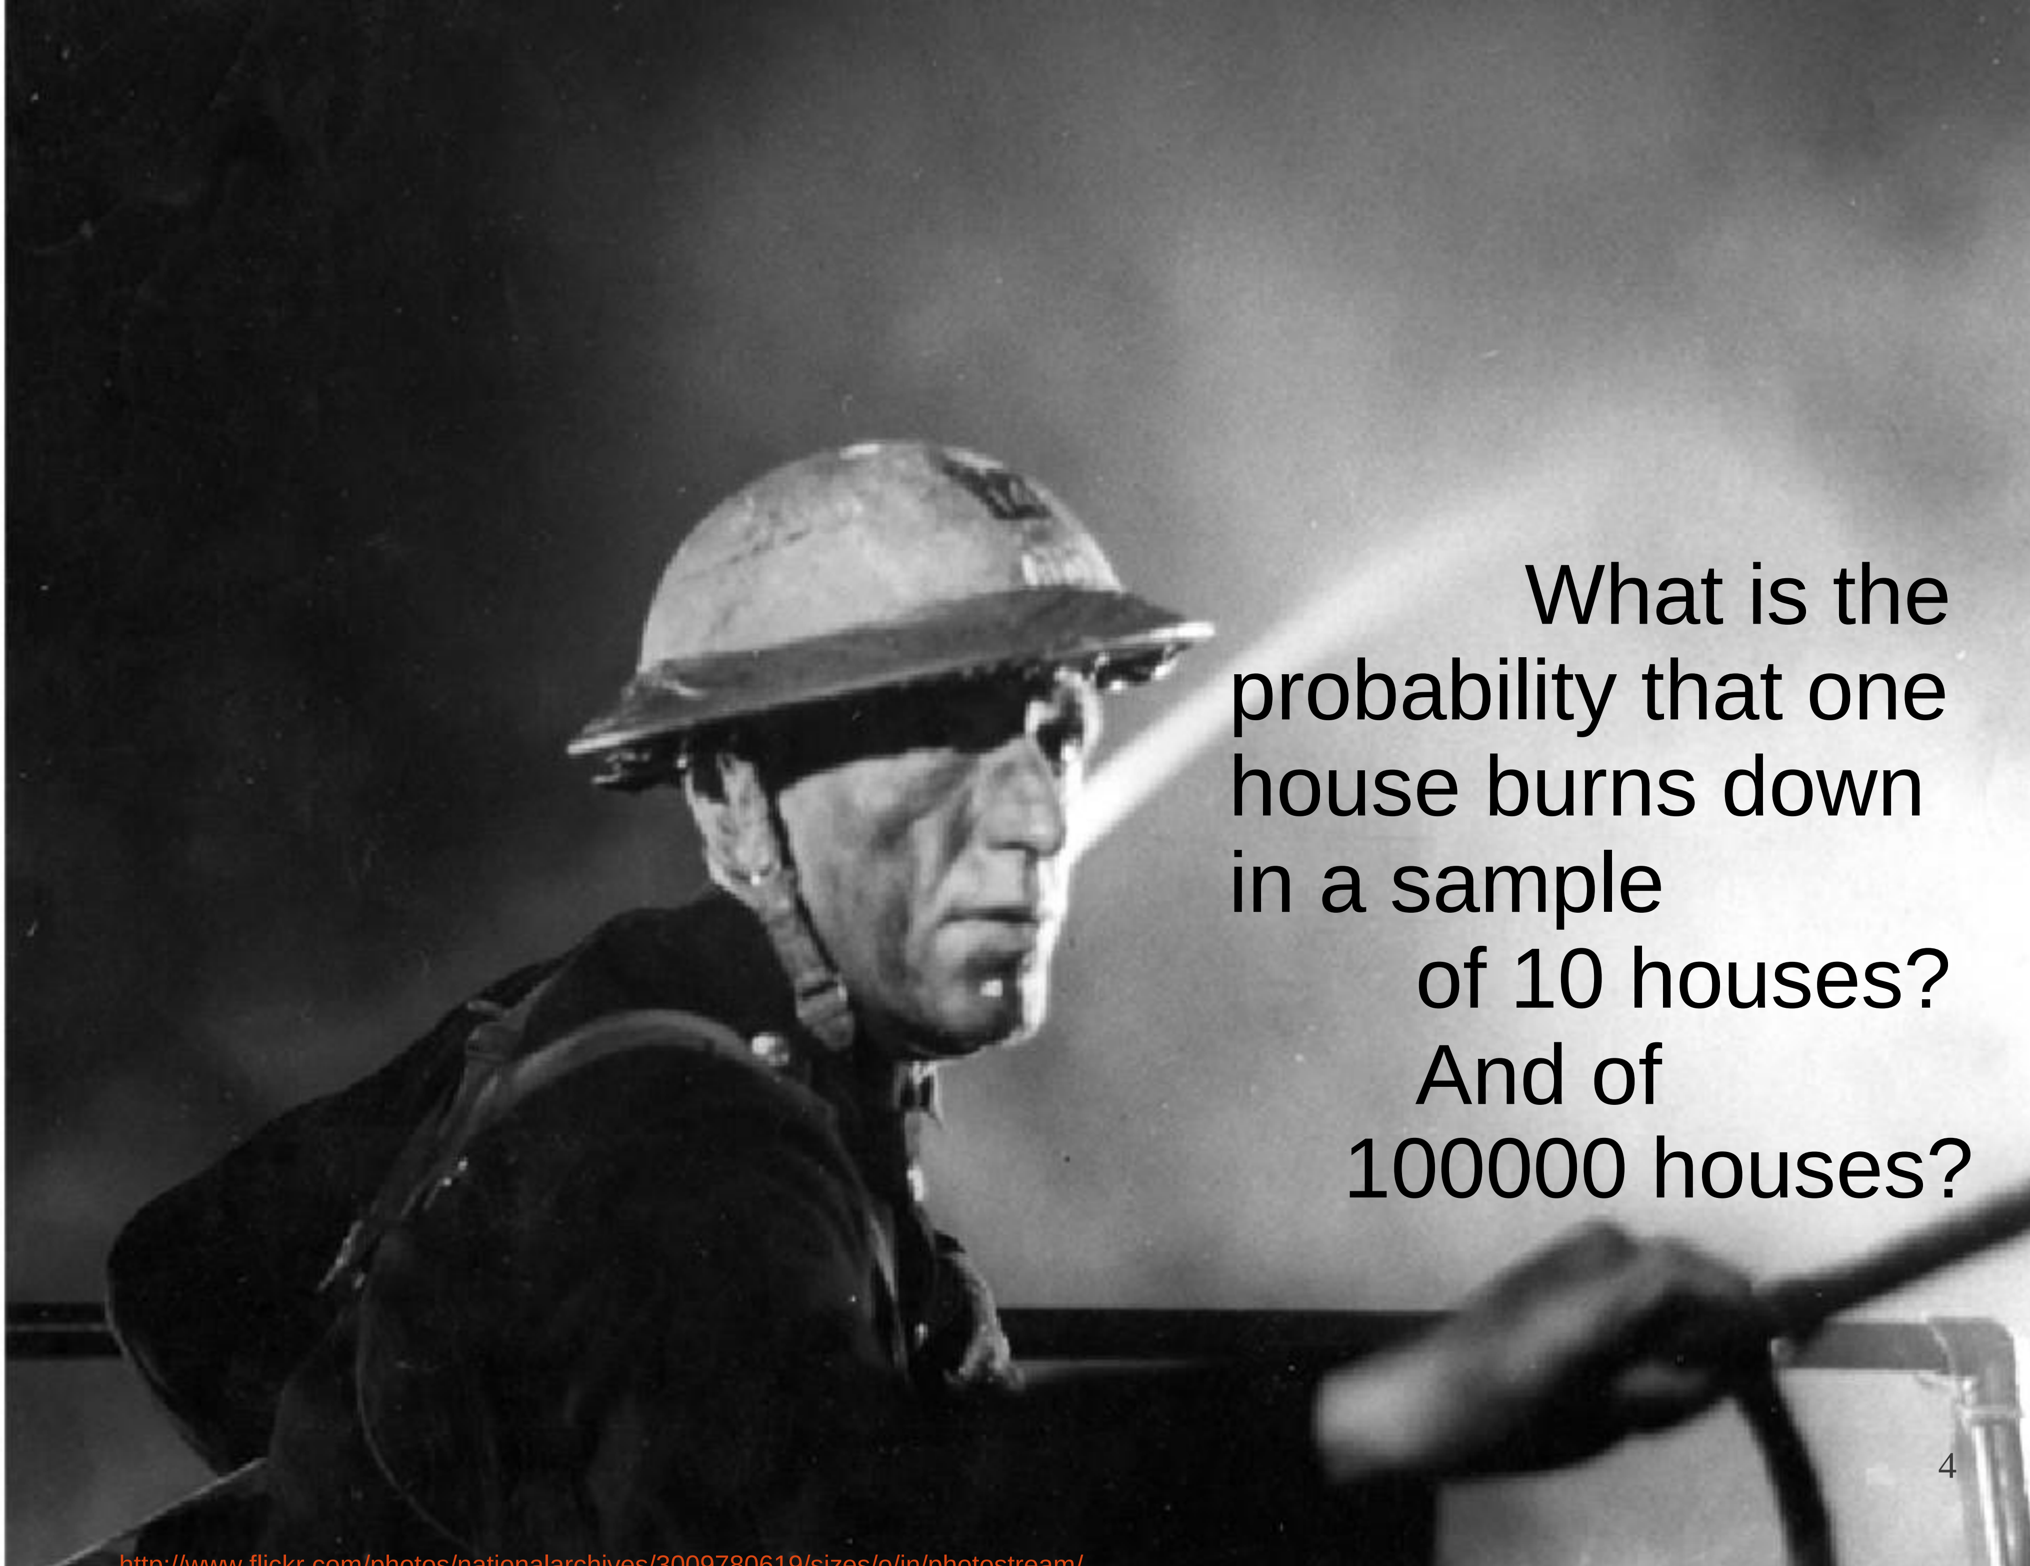

What is the probability that one house burns down in a sample of 10 houses? And of 100000 houses?

http://www.flickr.com/photos/nationalarchives/3009780619/sizes/o/in/photostream/

3: If heads, win R200 000, if tails, lose R100 000

With a small sample, Gamble 3 is not attractive
With a large sample, it starts to be.

We see **risk pooling** in:

Joint ownership of business ventures, racehorse syndicates, partnerships, joint stock companies

And of course:

Insurance markets. Why ?

Insurance: evidence of risk aversion

We have shown: most people would rather take a

small unfair gamble - paying premiums,
over a
large fair one - going without insurance.

Why can't insurance companies offer policies which are fair gambles?

What is the most people will pay for insurance?

Figure 5.7: The Reservation Price for Insurance

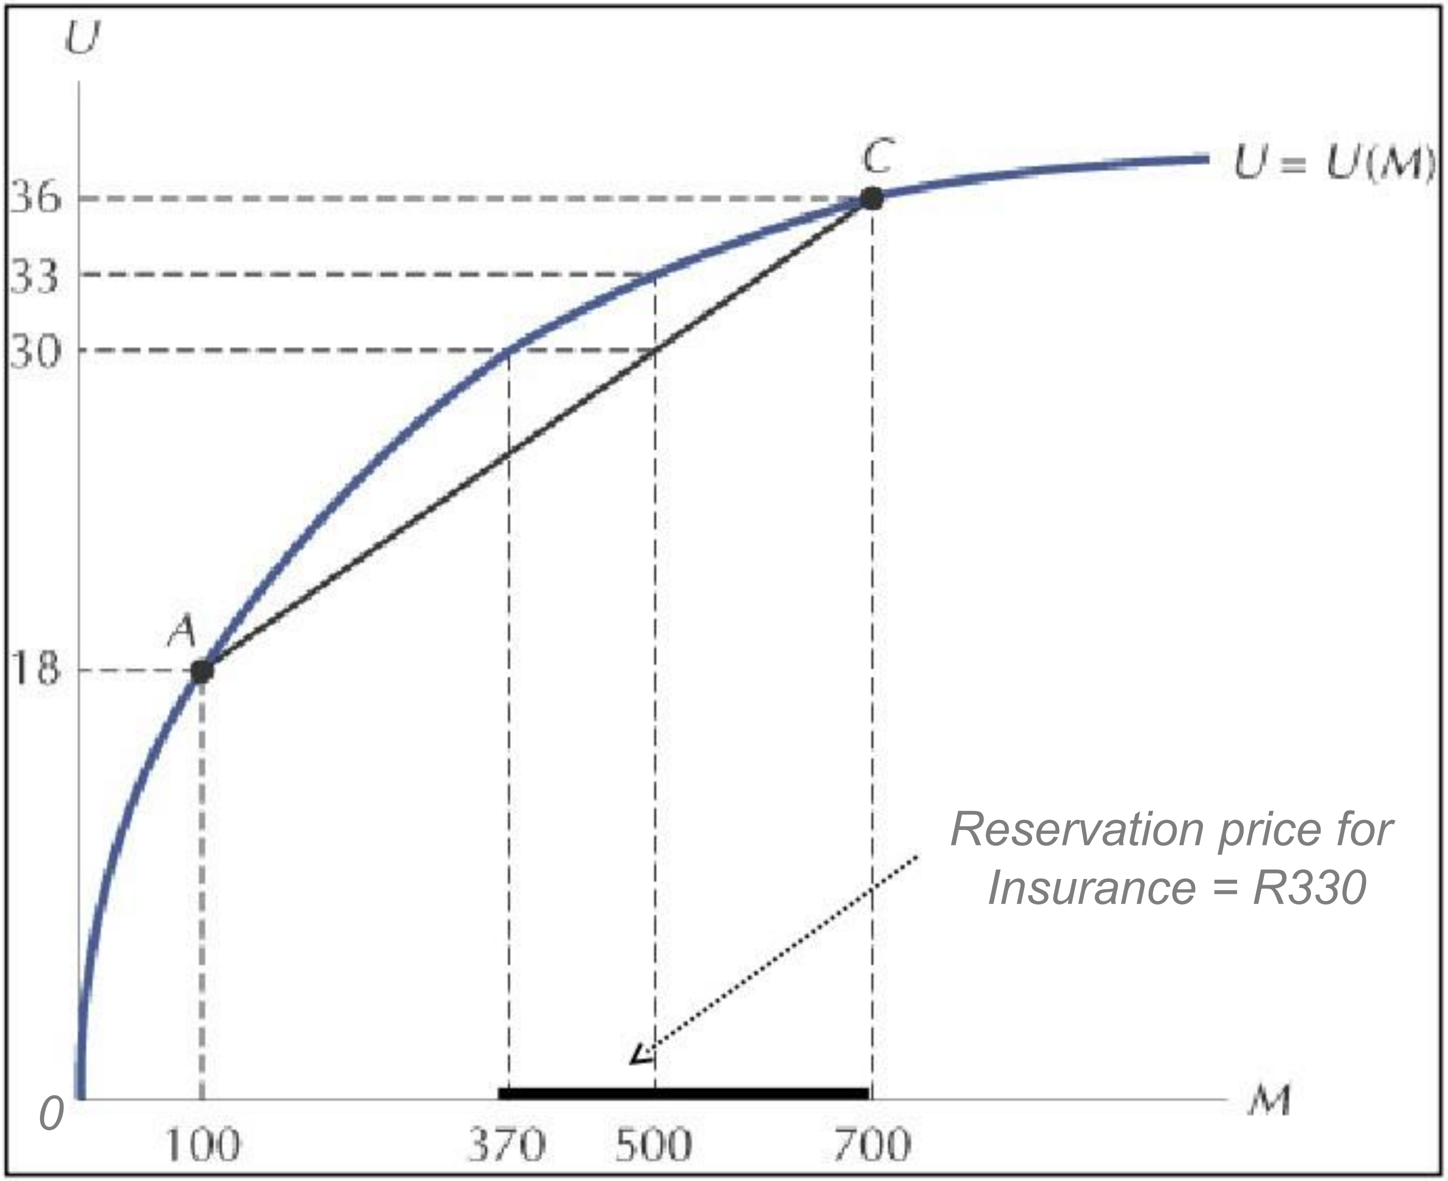

Nosipho has initial wealth R700, and faces a loss of R600 with probability 1/3

EU = 1/3U(100) + 2/3U(700) = 30
(See graph on next slide and P183)

This EU lies on the chord joining A and C

With no insurance,
expected wealth = 1/3(700-600) + 2/3(700) = 500

How much would she pay for a policy which completely covers the loss?

Nosipho's EU = 30 without insurance

What is the equivalent wealth which would give Nosipho U of 30 for certain? It is 370,

i.e. paying 700-370 = 330 to avoid the gamble

If she pays 330, she is **indifferent** between the gamble and taking insurance -
this is the maximum she is willing to pay, and is her **reservation price**

370 is the **certainty equivalent** value of the gamble

http://upload.wikimedia.org/wikipedia/commons/4/4d/Joshua_Tree_-_Love_car.jpg

Adverse Selection

A transaction takes place, and the less desirable trading partners are more likely to volunteer to exchange

Market for lemons with second hand cars, insurance markets

Arises from **lack of information**

Are consumers of insurance representative of the general population?

2 types of cars

Good ones - worth R100000
Bad ones - worth R50000

Equal numbers of good and bad at first

Buyers can't tell which is which

The average price buyers will pay is R75000

At that price, who agrees to sell?

http://www.flickr.com/photos/sugarhiccuphiccup/4808604692/sizes/l/in/photostream/

Adverse Selection

Can insurance companies identify the **riskiest drivers** and charge them more?

Premiums will generally reflect the **average risk,** which will make them attractive for bad drivers, and unattractive for good ones.

Good drivers then **self insure**, leaving more bad drivers in the market, which drives premiums up, which drives more good drivers to self insure, etc,

Moral Hazard

Moral hazard is when we take risks, because we know we're not going to pay for the consequences

Incentives that lead people to file fraudulent claims, or be negligent in their care of insured goods...

the claimant and the insurance company **do not have the same information** - i.e. Outsurance don't know I drive like a maniac, or that I leave my phone in my car etc etc.

Insurance upgrades, anyone?

Why does moral hazard arise?

2 reasons

Asymmetric information, Poor Mechanism Design

“we see a lot of claims on windscreens, because the low excess on windscreen damage creates a **moral hazard** - people don't take care not to park under coconut trees or drive too close behind lorries”

http://www.vectorstock.com/i/composite/36,64/443664/suspicious-man-vector.jpg

Statistical Discrimination

Arises from a **lack of information**

Treat people a certain way **based on observable characteristics**, like group identity

Older people charged more for life insurance, high car insurance premiums for young men under 25, or in different areas - town vs Mowbray

Is this okay? Is it legal? Is it understandable?

Statistical Discrimination

What if insurance companies abandoned statistical discrimination?

Genuinely low risk people would move to companies which retained the old structure, and high risk people from other companies would move over to take advantage of the new rates

The company would be left with **only high risk individuals**, and would have to charge more for everyone

Can we force all companies to abandon statistical discrimination?

If so, what would happen?

?

When to self insure?

Insurance premiums must cover not only losses, but also **admin costs**, and the costs of **adverse selection** and **moral hazard**

For small losses, rather self insure, or pick a higher excess

Premiums are lower for policies with higher excess values - why?

Small vs Large Losses

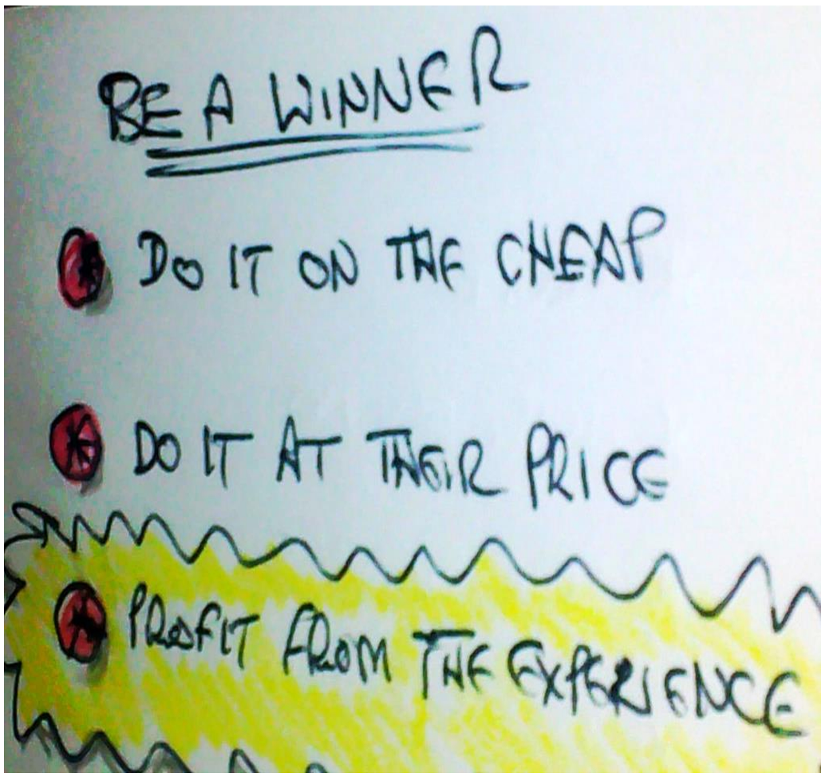

the winners curse: In most auctions, the winning bid generally exceeds the actual value of the item

The size of the winning bid depends on the behaviour of the **highest bidder**, not the average bidder.

Thus to earn a profit, auction off your item?

An **unbiased** estimate of a variable, on average, will be equal to the **true value**

- e.g. temperature forecasts are usually too high/low, but the long run average values track the actual values almost perfectly

Each bidder's estimate will be **too high** sometimes, and **too low** sometimes, but on average it will equal the true value of the item

The problem is that an auction is won by the bidder whose estimate is too high by the largest margin

A **fully rational bidder** will take into account the fact that the winning bid tends to be too high, and can adjust his/her bid downwards

The question is, by how much?

The expected value of the winning bid will be higher, the higher the number of bidders

So the more bidders there are, the more we should **adjust our estimate downwards**

Do bidders really adjust their bids downwards to avoid the winners' curse?

Those who don't will lose money on each auction, and will go bankrupt and **exit the market**

Fine, for a market of oil tenders, but what about you and I at the local antique auction? Can we avoid the winners' curse?