

Utilities and Information (Ch. 16)

Strong Dominance

	Good Chef	Bad Chef
Hamburger	 	 
No main	 	 

Weak Dominance

	Good Chef	Bad Chef
Hamburger (if good chef) No Main (otherwise)	 	 
No main	 	 

Announcements

HW3 posted (due 3/31)

Utility

Last time hopefully we motivated why utility is a good and expressive form of measurement (though remember: utility \neq money)

Using this as a basis, we will look at more at how you can reason efficiently

And also at how extra information corresponds to added utility

Multivariate Utility

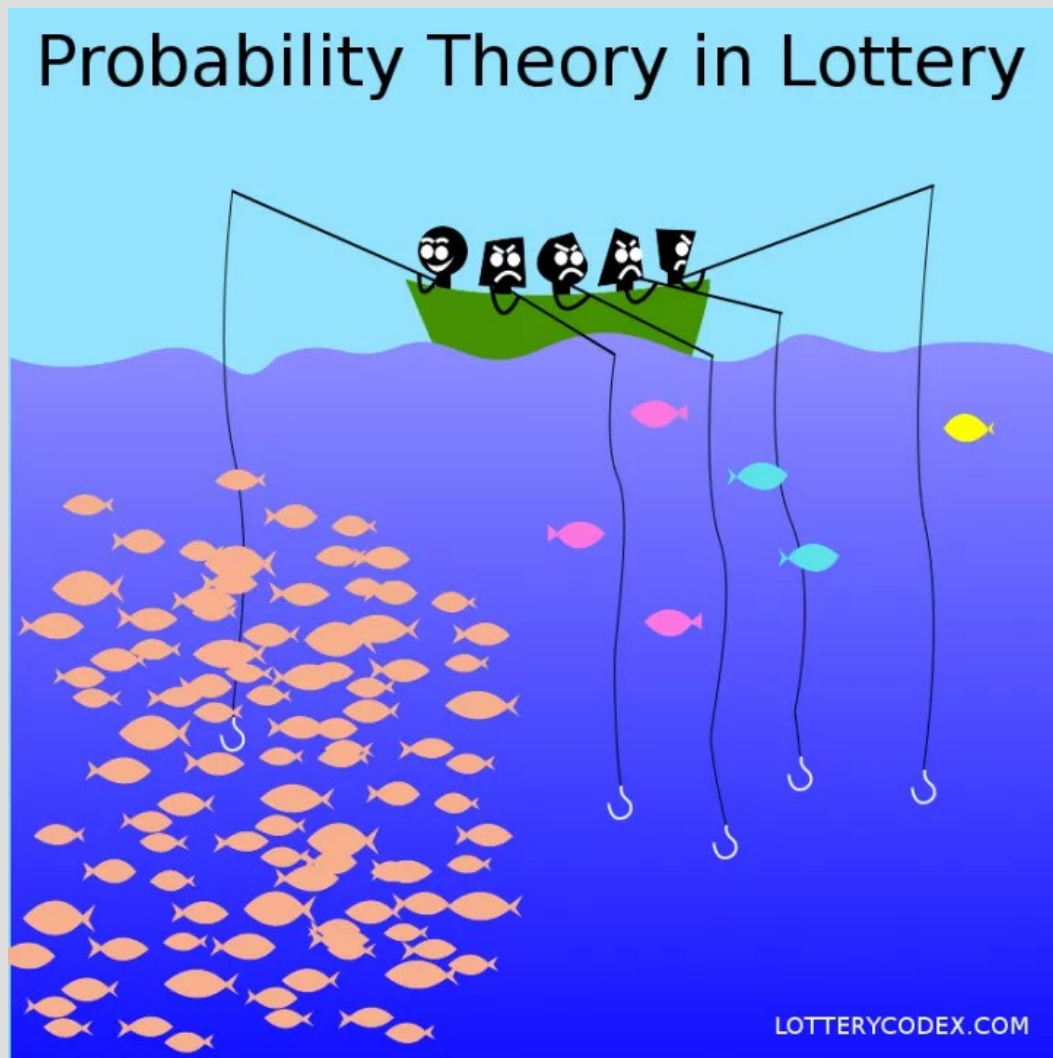
So far we focused on single-variable utilities

$$U(\text{seat}=\text{left}) = 8$$

$$U(\text{seat}=\text{middle}) = 6$$

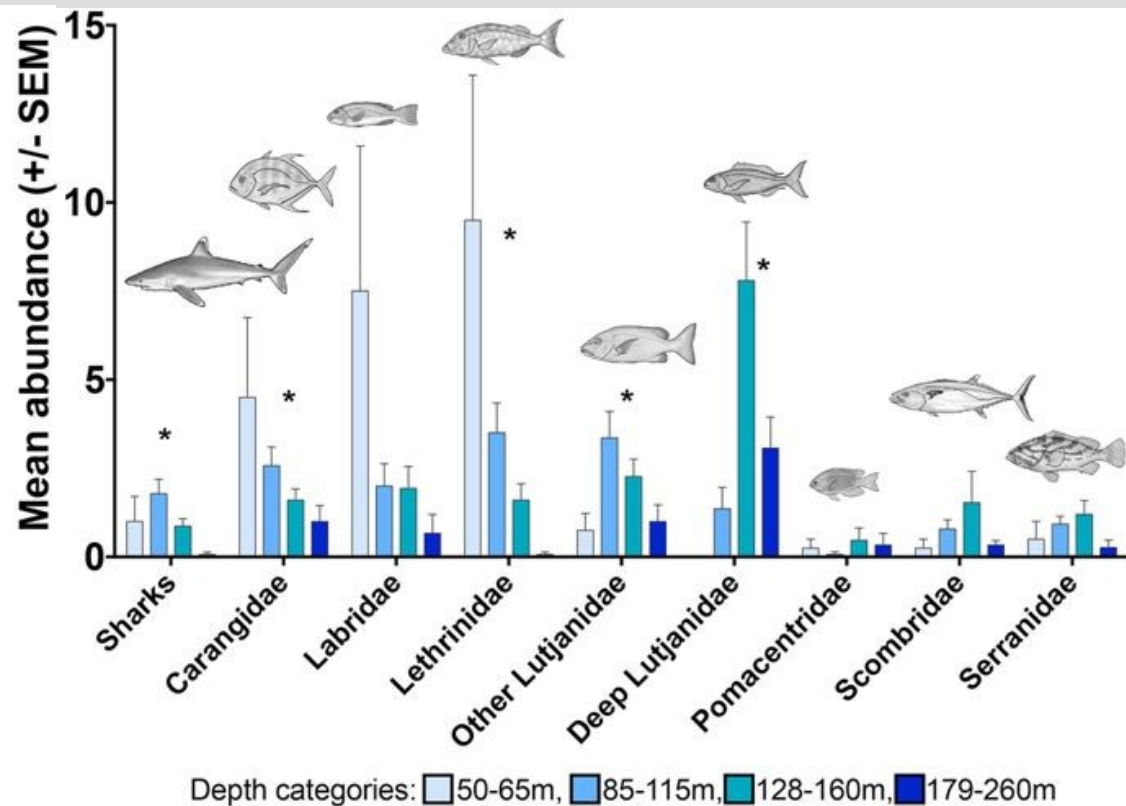
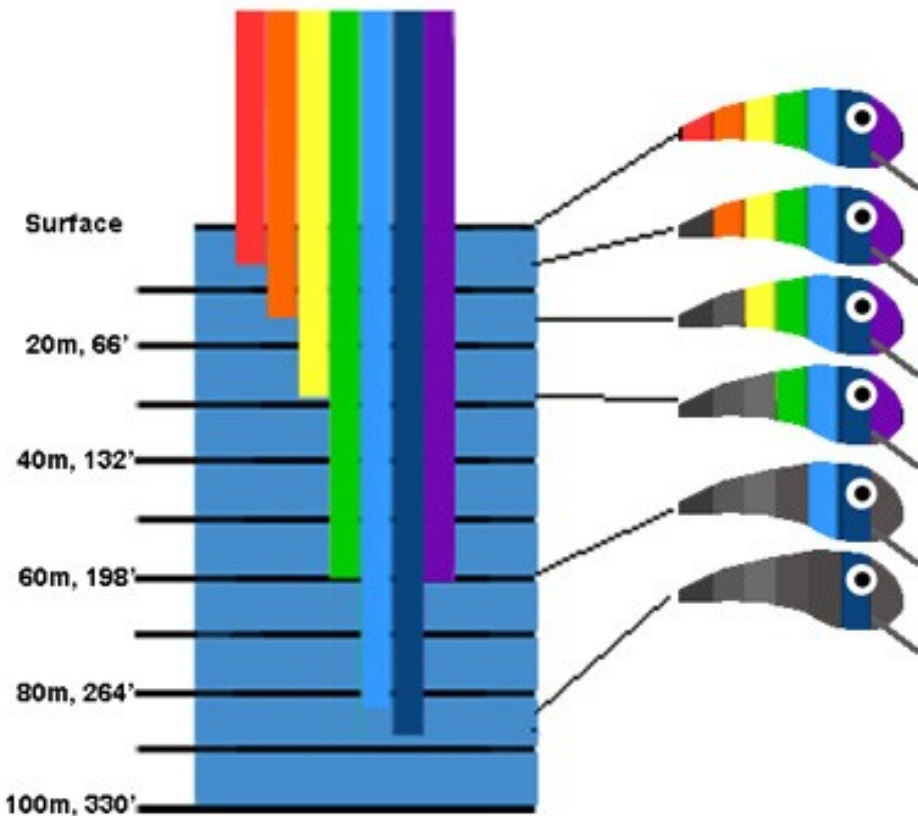
$$U(\text{seat}=\text{right}) = 2$$

But we could expand it to more than just the “seat” variable



Multivariate Utility

We could expand our “fishing” example to have (position, depth, lureColor)



Multivariate Utility

Now we would need to specify a utility for every combination of these variables....

$$U(\text{seat}=\text{left}, \text{depth}=5\text{m}, \text{Lure}=\text{RGB}) = 8$$

$$U(\text{seat}=\text{left}, \text{depth}=5\text{m}, \text{Lure}=\text{GB}) = 6$$

...

Depending on how many variables and values per variable you have... this can be quite exponential, so we have to find a better way

Multivariate Utility

Let's go down to 2 variables (seat, depth):

$$U(\text{seat}=\text{left}, \text{depth}=5\text{m}) = 8$$

$$U(\text{seat}=\text{left}, \text{depth}=10\text{m}) = 5$$

$$U(\text{seat}=\text{middle}, \text{depth}=5\text{m}) = 6$$

$$U(\text{seat}=\text{middle}, \text{depth}=10\text{m}) = 7$$

$$U(\text{seat}=\text{right}, \text{depth}=5\text{m}) = 2$$

$$U(\text{seat}=\text{right}, \text{depth}=10\text{m}) = 6$$

“seat=right” is worse for all depth values than “seat=middle”, called strictly dominated

Multivariate Utility

What does being strictly dominated tell us?

How important are the utilities?
(What other assumptions could you make?)

Multivariate Utility

If something is strictly dominated, we can ignore the choice in our decision making (as it is always terrible)

You can find dominance even if you don't know the actual utility values, as you quite often know if a factor is “good” or “bad”

Let's change examples from fishing to jobs with properties (fun, pay)

Multivariate Utility

Obviously you want a job that pays well and is also exciting, but your job choices are limited to a few options (i.e. actions)

McDonalds = (fun=-3, pay=1)

Teacher = (fun=4, pay=3)

Banker = (fun=0, pay=6)

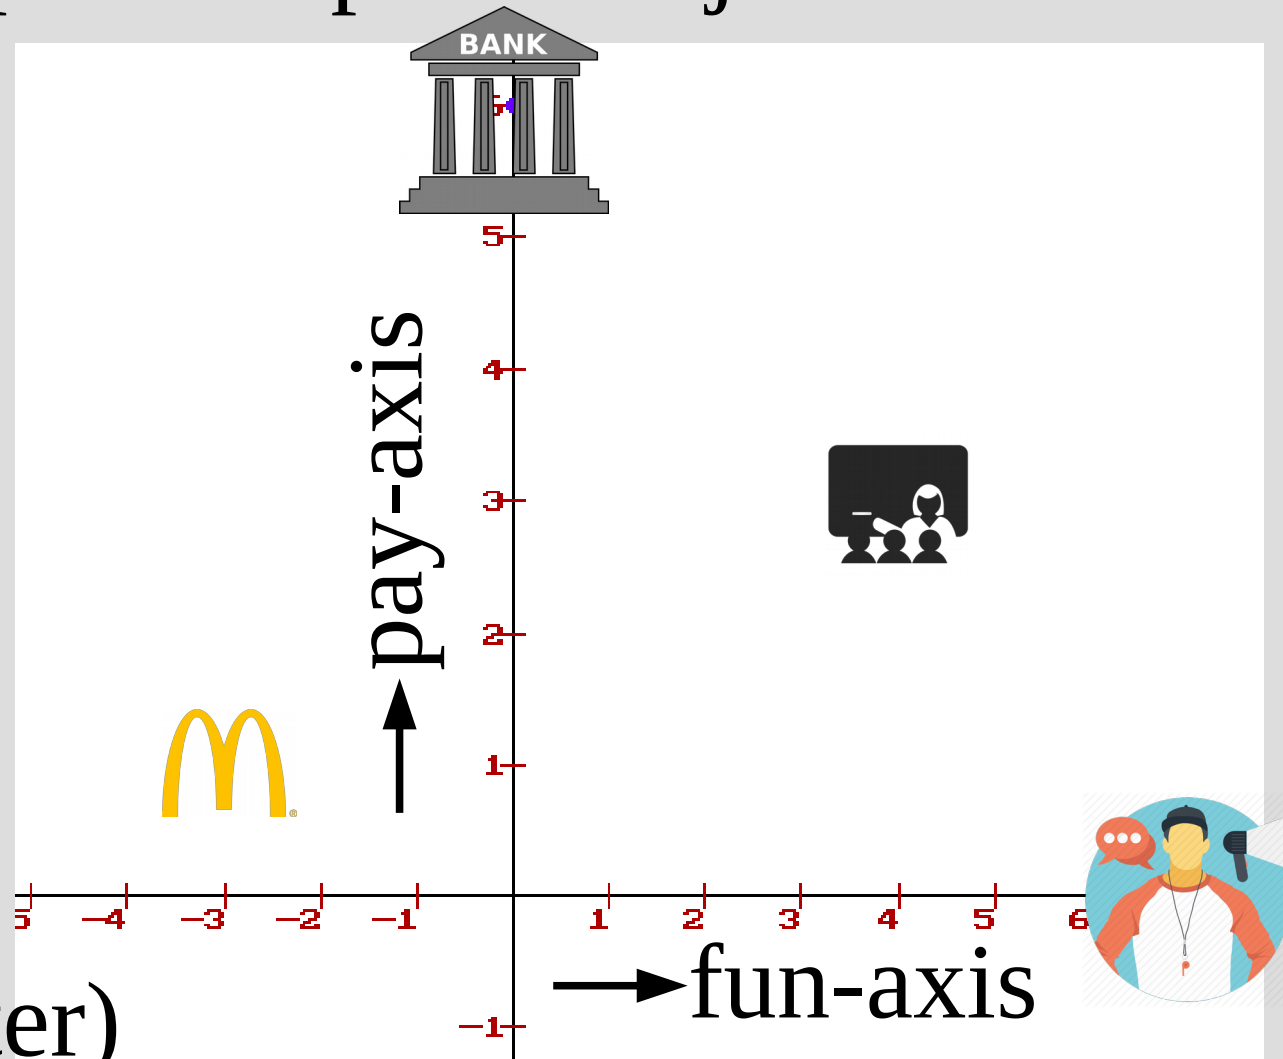
Volunteer = (fun=7, pay=0)

Multivariate Utility

You could then plot the possible jobs:

Although you do not have a utility for jobs

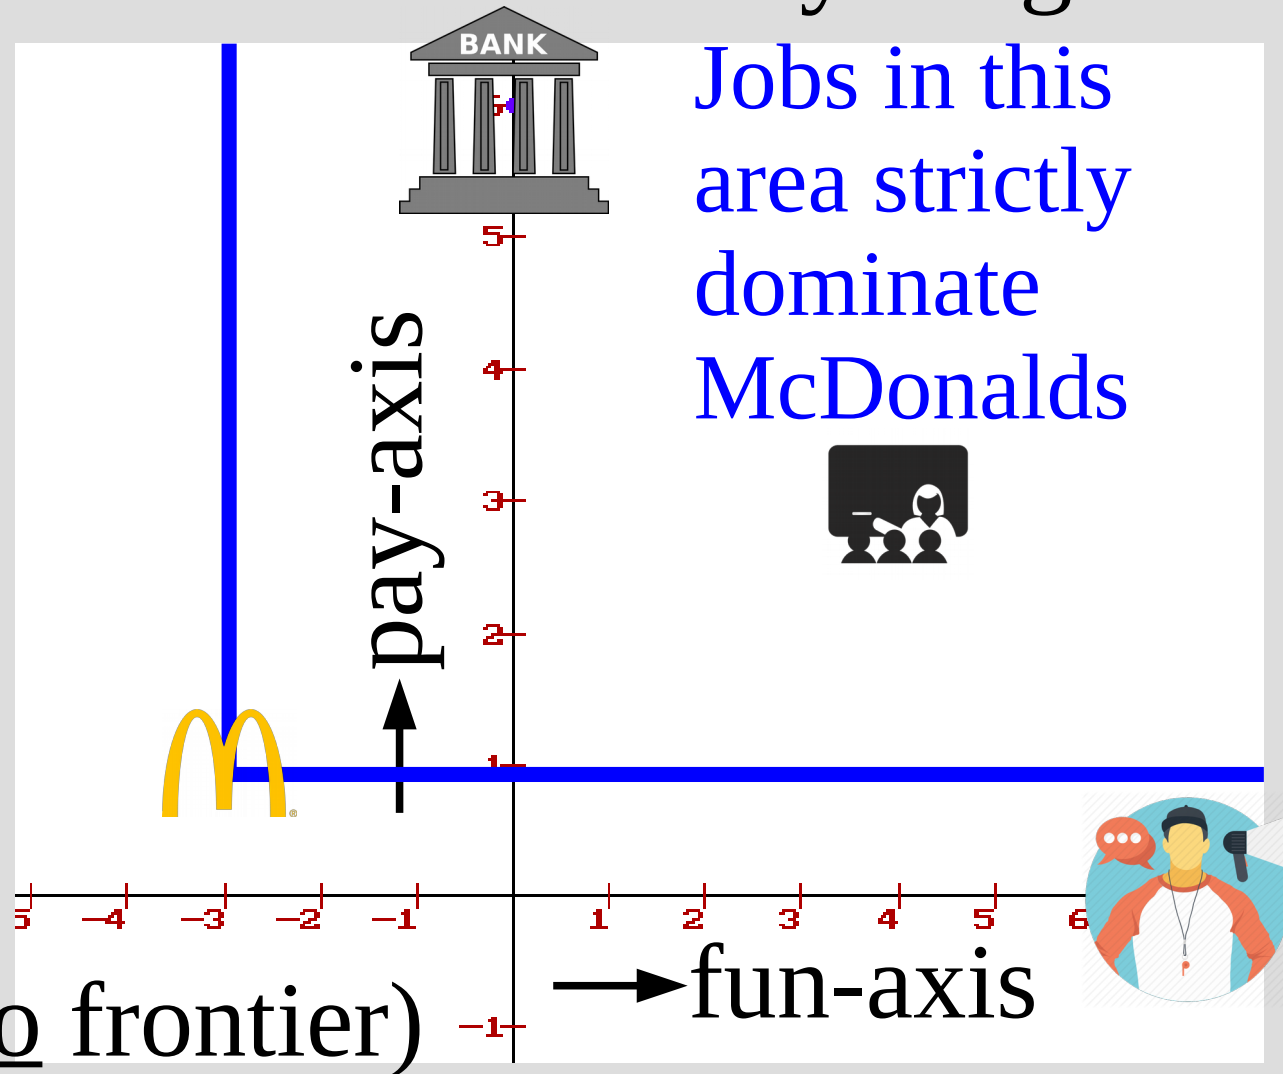
Fun and pay are monotonically increasing (i.e. more is better)



Multivariate Utility

So you can figure dominance as anything that has both better (or equal) pay and fun

Thus there is no reason to consider McDonalds (related to Pareto frontier)



Stochastic Utility

However, not all parts might be fixed (i.e. uncertainty)

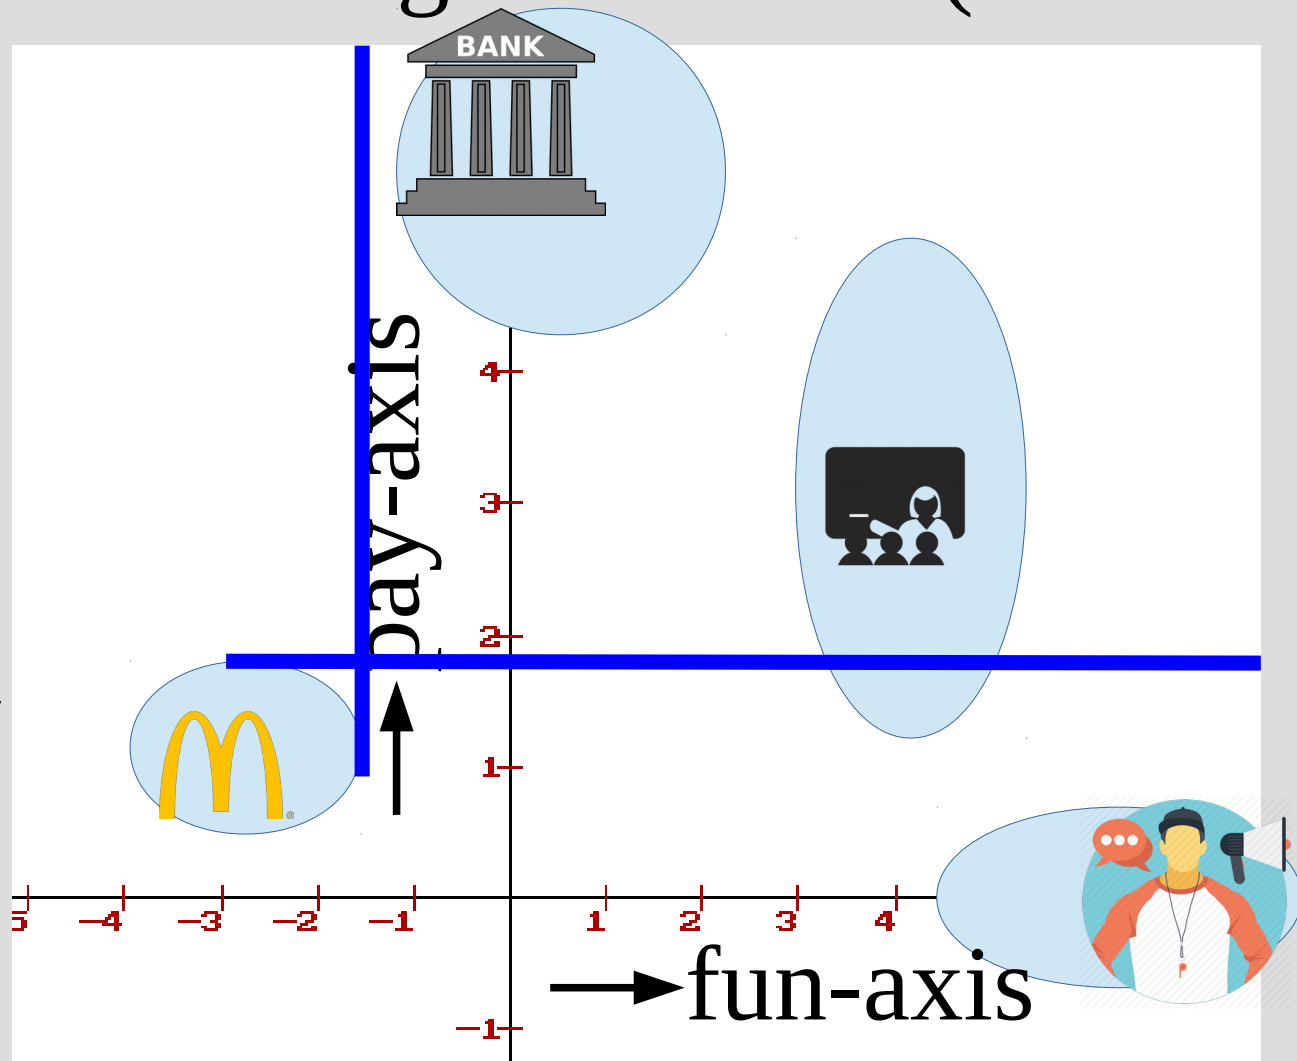
For example, you might want to compute the pay over a couple years, but you don't know when/if you will get promoted

Sometimes you might be stuck doing a boring job (flipping the burgers?) or less boring parts (taking orders?)

Stochastic Utility

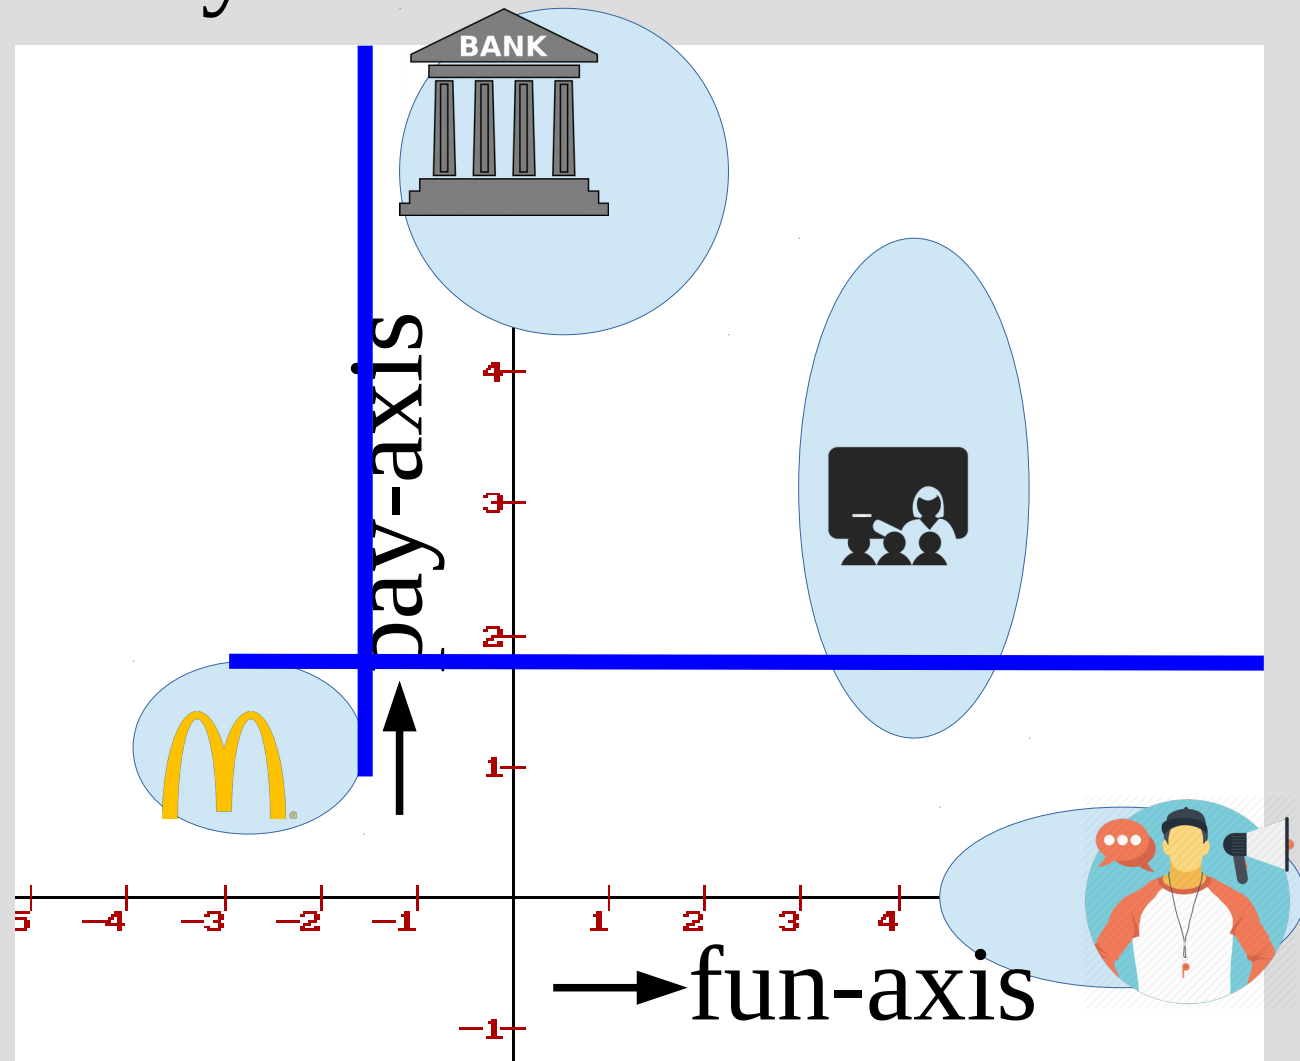
Thus we could have a range of values (with a distribution)

We can still have strict dominance if all of one area is up&right of another option (i.e. banking)



Stochastic Utility

Is there any other way we could define “dominance” when we have probabilities and/or distributions?

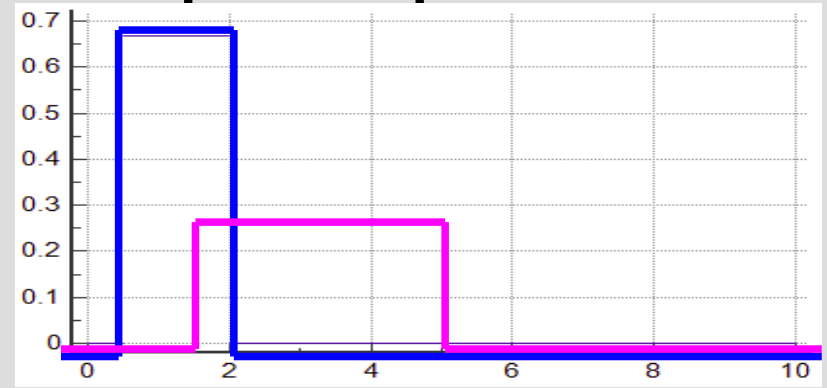


Stochastic Utility

You can still have dominance even if you do overlap though, but you need one option to always be better still

Consider just the pay side and assume both **McDonalds** and **Teaching** is uniformly distributed between $[0.5 \text{ and } 2]$ and $[1.5 \text{ and } 5]$

We would call this stochastically dominant



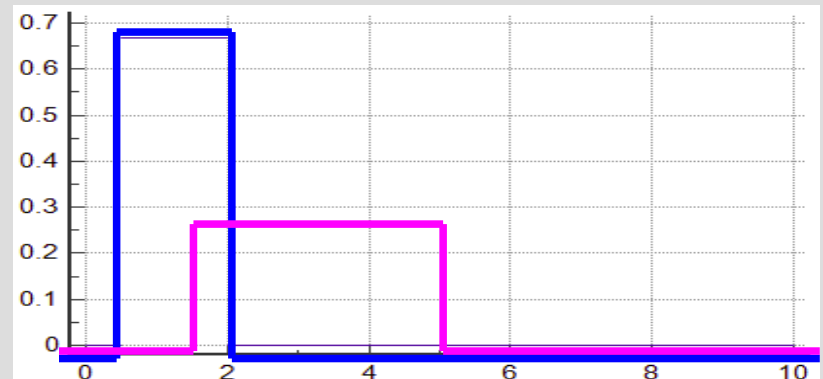
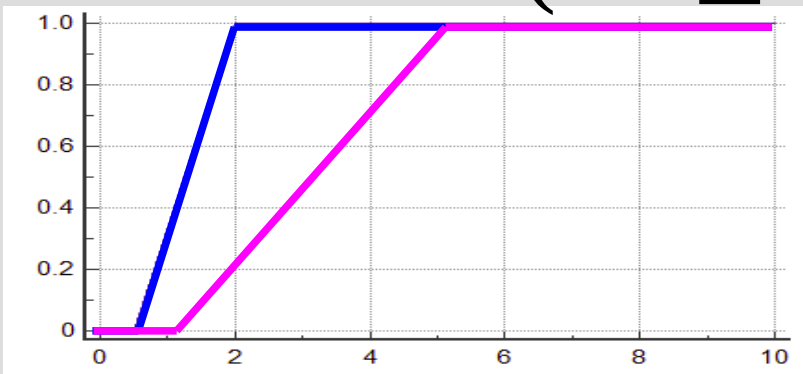
Stochastic Utility

Specifically, we need more area under the “worse” curve at all times to be stochastically dominant (from left to right, so more small):

$$\forall x \int_{-\infty}^x Mc(n)dn \geq \int_{-\infty}^x Teach(n)dn$$

distributions for these

These integrals can be visualized as ($Mc \geq Teach$):



Stochastic Utility

Note, this is different than comparing just the expected utility or if you knew the value

In these cases you would actually need to know the value of money to figure out if:

$$p * U(\$0.5) + (1-p) * U(\$2) \geq U(\$1.7)$$

You can use dominance (either type) to eliminate “bad” options, but rarely does this leave you with only 1 choice (so more work)

Utility Simplifications

Fortunately, we can avoid specifying an exponential number of utilities (sometimes)

We define preference independence as (assuming 3 variables/attributes):

Variables x and y are preferentially ind. if (x_1, y_1) preferred over (x_2, y_1) means for all y : (x_1, y) preferred over (x_2, y)

Utility Simplifications

This may seem like a strong requirement, but it is normally true (if the things you are measuring (i.e. axis) are independent)

For example:

Job1=(fun=2, pay=3), Job2=(fun=1, pay=3)

... You prefer Job1 over Job2 (more fun)

But this is true for any pay amount: $(A \succ B)$

JobA=(fun=2, pay=x), JobB=(fun=1, pay=x)

Utility Simplifications

We can expand this definition to sets of variables as well:

JobA=(fun=2,dist=3,pay=5,time=8)

JobB=(fun=4,dist=1,pay=5,time=8)

JobA preferred to JobB (more fun, closer)

We would say the set {fun,dist} is preferentially independent of {pay,time} if:
(fun=2,dist=3,pay=x,time=y) preferred to
(fun=4,dist=1,pay=x,time=y) for any x,y

Utility Simplifications

If in a subset of variables, A , all are preference independent from each other (over all subsets of A), then mutually preference independent

Then, you can actually say their utility is additive! (very non-exponential work)

$$U(\text{fun}=w, \text{pay}=x, \text{dist}=y, \text{time}=z) = \\ a * U(\text{fun}=w) + b * U(\text{pay}=x) + c * U(\text{dist}=y) \\ + d * U(\text{time}=z)$$

Utility Simplifications

So far we have been assuming the variables have “fixed” values (i.e. no uncertainty)

Things become a bit more problematic if we involve probabilities...

We can still have “mutual independence”, though this time we call it mutually utility independent (not mutual preference ind.)

Utility Simplifications

Thankfully, utility independence is similar to preference independence

A random variable x is independent to a (random or normal) variable y if:

(x_1, y_1) preferred over (x_2, y_1) means
for all y : (x_1, y) preferred over (x_2, y)

Utility Simplifications

Going back to the job example, say you have a random variables x_1, x_2 :

$$x_1 = [(0.5, \$2), (0.5, \$4)]$$

$$x_2 = [(0.2, \$0), (0.8, \$6)]$$

Assume you prefer JobA over JobB in:

$$\text{JobA}(\text{fun}=2, \text{pay}=x_1), \text{JobB}(\text{fun}=2, \text{pay}=x_2)$$

If utility independent, true for any “fun” z :

$$\text{JobA}(\text{fun}=z, \text{pay}=x_1) \succ \text{JobB}(\text{fun}=z, \text{pay}=x_2)$$

Utility Simplifications

Unfortunately in the probability case, this independence does not lead to simple additive

Instead, utility independence for set {a,b,c}:

$$\begin{aligned} U(a, b, c) = & k_a \cdot U_a(a) + k_b \cdot U_b(b) + k_c \cdot U_c(c) \\ & + k_a \cdot k_b \cdot U_a(a) \cdot U_b(b) + k_a \cdot k_c \cdot U_a(a) \cdot U_c(c) + k_b \cdot k_c \cdot U_b(b) \cdot U_c(c) \\ & + k_a \cdot k_b \cdot k_c \cdot U_a(a) \cdot U_b(b) \cdot U_c(c) \end{aligned}$$

Or in general: different utility functions as one is for “fun” and another for “pay”

$$1 + U(x_1, x_2, \dots, x_n) = \prod_{i=1}^n \left(1 + k_i \cdot U_i(x_i) \right)$$

why it's also called “multiplicative” independence

Utility Simplifications

You can get an additive utility function even with random variables, but you need another fact to hold true:

We need to be able to treat the combination of variables as random variables as well, like:

$$\text{OptionA} = [(0.5, x_1), (0.5, y_1)]$$

$$\text{OptionB} = [(0.5, y_1), (0.5, y_2)]$$

Utility Simplifications

Going back to the job example, assume 2 “pay” variables (x_1, x_2) and 2 “fun” (y_1, y_2)

If we work at McDonalds, 50% which position
 $[(0.5, (\text{pay}=x_1, \text{fun}=y_1)), (0.5, (\text{pay}=x_2, \text{fun}=y_2))]$

Suppose another job somewhere else has:

$[(0.5, (\text{pay}=x_1, \text{fun}=y_2)), (0.5, (\text{pay}=x_2, \text{fun}=y_1))]$

swap



If these jobs are “equal”, then also additive

Utility Simplifications

There are some more nitty-gritty cases, like when x is utility independent of y ...
but you could have y not be independent of x

You could see this in a pay “by the hour” job
“pay” would be independent of “time” (more pay always better with same time on job)...

...but “time” not independent with “pay”
(might want to take more/less time on job)

Utility Simplifications

Giving actual numbers to this:

Assume you want to make over \$20,000 a year
(above poverty line... for a family of 3)

So you are very unhappy (utility) if you get below this amount, but over this amount your happiness only increases slowly

Job1=(\$15k, 40 hr/w), Job2=(\$15k, 80 hr/w)

Job3=(\$25k, 40 hr/w), Job4=(\$25k, 80 hr/w)

Utility Simplifications

Job1=(\$15k, 40 hr/w), Job2=(\$15k, 80 hr/w)
Job3=(\$25k, 40 hr/w), Job4=(\$25k, 80 hr/w)

Up and down, more pay is always better

But if you didn't have Job3 as an option,
the best might be Job4 (sell your soul...)

This means you would prefer Job4 over Job1,
which has a higher \$/hr pay (or not work 80hr)

Utility Simplifications

Job1=(\$15k, 40 hr/w), Job2=(\$15k, 80 hr/w)
~~Job3=(\$25k, 40 hr/w), Job4=(\$25k, 80 hr/w)~~

Up and down, more pay is always better

But if you didn't have Job3 as an option,
the best might be Job4 (sell your soul...)

This means you would prefer Job4 over Job1,
which has a higher \$/hr pay (or not work 80hr)

Utility of Information

Now that we can compute values (in a hopefully non-exponential way) we can also measure how “useful” information is

Remember from last time expected utility is:

$$EU(\alpha|e) = \max_a \sum_{s'} P(\text{Result}(a) = s' | a, e) \cdot U(s')$$

alpha is argmax a (the a which is the best)

probability taking action a makes you end up in state s' when you know e

We can now see the benefit of different information (via utility)

Utility of Information

Let's look at a (maybe) relevant example:

You getting ready to take a test in a class, you can either study for the test or not:

Study: 90% get state: (class=pass, fun=no)
10% get state: (class=fail, fun=no)

Play: 50% get state: (class=pass, fun=yes)
50% get state: (class=fail, fun=yes)

Utility of Information

Assume we can use an additive utility function

$$U(\text{class}, \text{fun}) = 4 * \text{class} + \text{fun}$$

$$(1=\text{true}, 0=\text{false} \dots U(\text{class}=\text{pass}, \text{fun}=\text{no}) = 4)$$

Then we could find the expected value of actions (as they are random variables):

$$EU(\text{study}) = 0.9 * 4 + 0.1 * 0 = 3.6$$

$$EU(\text{play}) = 0.5 * 5 + 0.5 * 1 = 3$$

... so you should study

Utility of Information

Let's say someone offers you the answer key for the test in exchange for money... then:

Study(with ans): 100% (class=pass, fun=no)

Play(with ans): 100% (class=pass, fun=yes)

$EU(\text{study} \mid \text{ans}) = 4$

$EU(\text{play} \mid \text{ans}) = 5$

So in this case you could just “play”

Utility of Information

The question is, how much money would you (rationally) pay for the answers?

Best action without answers: $EU(\text{study}) = 3.6$

Best action with answers: $EU(\text{play} \mid \text{ans}) = 5$

So you should be willing to pay 1.4 utility worth of money to get the answer key (this is the “value” or “utility” of the info)

please don't actually try to buy answer keys though...

Utility of Information

We can actually account for the case when the bought answers are “outdated”

Say 30% of the time, the answer key is incorrect (use original outcomes)

Then we could compute this as a more complex random variable:

$$E[\text{study\&maybeAns}] = [(0.3, [(0.9, 4), (0.1, 0)]) \\ , (0.7, [(1.0, 4)])]$$

Utility of Information

So, we would compute:

$$E[\text{study\&maybeAns}] = 0.3 * 3.6 + 0.7 * 4 = 3.88$$

$$E[\text{play\&maybeAns}] = 0.3 * 3 + 0.7 * 5 = 4.4$$

So even with the faulty answers you should choose “play” but you are only willing to offer (4.4-3.6) or 0.8 dollars worth of utility

It should make sense that this is down from 1.4 dollars with accurate answers

Utility of Information

If the “answers” to the test were wrong 80% of the time... you would compute

$$E[\text{study\&maybeMaybe}] = 0.8 * 3.6 + 0.2 * 4 = 3.68$$

$$E[\text{play\&maybeMaybe}] = 0.8 * 3 + 0.2 * 5 = 3.4$$

Here you would now want to study, as the “answers” are very unreliable

... but you would still want to buy them, just for 0.08 utilities of money

Value Perfect Information

The book defines this calculation of how much information/evidence is worth as value of perfect information:

$$VPI_e(E_j) = \left(\sum_k P(E_j = e_{j,k} | e) \cdot EU(\alpha_{e_{j,k}} | e, E_j = e_{j,k}) \right) - EU(\alpha | e)$$

Value Perfect Information

The book defines this calculation of how much information/evidence is worth as value of perfect information:

$$VPI_e(E_j) = \left(\sum_k P(E_j = e_{j,k} | e) \cdot EU(\alpha_{e_{j,k}} | e, E_j = e_{j,k}) \right) - EU(\alpha | e)$$

sum over key correct/not

probability answer key correct

expected best result with (in)accurate key

expected best result without key

Value Perfect Information

“Perfect information” might be bad wording for our example as our information was not perfect (it was not correct sometimes)

You could convert it into an example that does have perfect information:

The answer key always exists and is correct but the thug your hire only has a 70% chance at successfully stealing the key (yikes!)

Value Perfect Information

The perfect information equation also incorporates the fact that you have some initial evidence as $EU(\alpha|e)$ is our “baseline”

We could work this into the example as $VPI_{e=\{\}}(\text{key})=70\%$ thug successful=1.4 (old calc)

If you then find out I own a dog, and will decrease the thug's success rate to 20%

$VPI_{e=\{\text{dog}\}}(\text{key})= 20\%$ thug steals =0.08 (old calc)

Value Perfect Information

What sort of properties would you want when you “evaluate” information?

(i.e. what should be true about VPI?)

Value Perfect Information

This formulation has some nice properties:

(1) VPI positive:

$$\forall e, x: VPI_e(x) \geq 0$$

(2) VPI order independent

$$\forall e, x, y: VPI_e(x, y) = VPI_e(x) + VPI_{e,x}(y) = VPI_e(y) + VPI_{e,y}(x)$$

(3) VPI is not additive (counter property?)

$$VPI_e(x, y) \neq VPI_e(x) + VPI_e(y)$$