Uncertainty

Noah Lyman

June 1, 2023

Introduction

- We've said previously that there are two ingredients to a utility maximization problem:
  1. Preferences
  2. Constraints
- Truly, there are three ingredients:
  1. Preferences
  2. Constraints
  3. **Information**
- To characterize optimal decisions, we need to know:
  1. What the agent likes
  2. What constraints the agent faces
  3. What information the agent has access to

Introduction

- Consider two investors deciding how best to allocate their budget over a stock and a bond
- To the best of Investor A's knowledge, the stock has a 20% chance of increasing in price
- Investor B has insider information, and knows there is a 95% chance the stock will increase in price
- Due to their differing information, investors A and B will probably purchase different amounts of the stock

Introduction

- Consider a different example in which two consumers decide whether or not to purchase health insurance
- Both consumer A and consumer B have a heart condition
- Consumer A is not aware of their heart condition, so does not anticipate needing medical care
- Consumer B knows of their heart condition, so anticipates a need for medical care soon
- Again, due to their differing information, it seems reasonable to think that consumer B is more likely to purchase insurance than consumer A

Introduction

- Up to this point, we've assumed that agents have perfect information
  - They have perfect knowledge of the state of economic environment
- It is more realistic to assume people have limited information
  - They face some *uncertainty*
- Examples:
  - Purchasing a stock: who knows if the price will go up or down?
  - Purchasing health insurance: who knows if I'll get sick or not?
  - Picking a college major: who knows if I'll be good at it or not?
- We'll begin this section by discussing how to model uncertainty mathematically

Lotteries

- We begin with the set X of possible events
  - The set of all things which can possibly happen
  - Referred to as the *sample space*
- A *lottery* ℓ is a probability distribution over X
  - Specifies the probability of each event occuring
- Example: flipping a coin
  1. $X = \{H, T\}$
  2. $\ell = (\frac{1}{2}, \frac{1}{2})$
- Example: rolling a die
  1. $X = \{1, 2, 3, 4, 5, 6\}$
  2. $\ell = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})$

Lotteries

$$\ell = (p_1, p_2, \ldots, p_N)$$

- Given a set $X = \{p_1, p_2, \ldots, p_N\}$ of N events, the corresponding lottery ℓ is a set of N probabilities
  - One for each event
- A few important notes about probabilities are in order:
  - Probabilities are always weakly positive:

$$p_i \geq 0 \quad \text{for all} \quad i = 1, \ldots, N$$

  - Probability distributions (i.e. lotteries) sum to 1:

$$\sum_{i=1}^{N} p_i = p_1 + p_2 + \ldots + p_N = 1$$

Interpreting Probabilities

- How do we interpret probabilities? This is the matter of philosophical debate, but broadly speaking there are two interpretations.

- First, probabilities can represent the **objective** likelihood of something happening
  - Ex: if rolling a fair die, there is objectively a $\frac{1}{6}$ chance of rolling a 1
  - **Frequentist interpretation**

- Alternatively, probabilities can represent one's **subjective** beliefs about the likelihood of something happening
  - Ex: A coach believes they'll win with .8 probability and lose with .2 probability
  - **Bayesian interpretation**

- No matter the interpretation, the mathematical treatment of probabilities remains the same

Expectations

- Let's consider a simple game:
  - Flip a coin
  - If heads, you give me \$1
  - If tails, I give you \$1
- On average, how much money will a make per turn?
- Below I plot a simulation of 20 repetitions of this game
  - Vertical axis is average winnings after *n* games

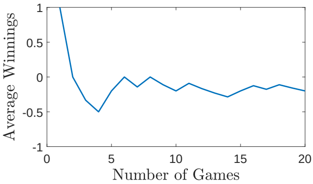

- What if we played 100 times?

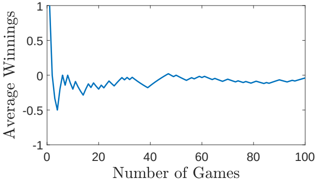

- What if we played 1000 times?

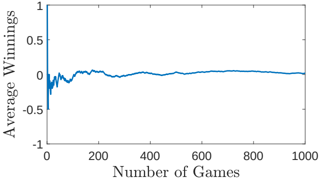

Expectations

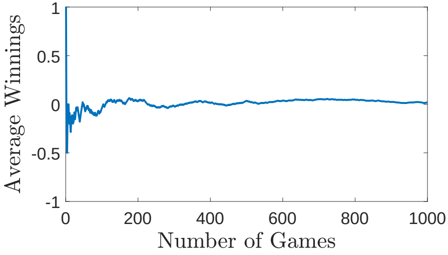

- Back to our question: On average, how much money will a make per turn?
- It seems like the correct answer is \$0, but where does this answer come from?

Expectations

- The answer lies in the concept of *expectations*
- Given a random variable x which can take values in $X = \{x_1, x_2, \ldots, x_N\}$
- The probabilities of each realization are specified by the lottery $\ell = (p_1, p_2, \ldots, p_N)$
- Then the expected value of x is given by:

$$\begin{aligned}\mathbb{E}[x] &= \sum_{i=1}^{N} p_i x_i \\ &= p_1 x_1 + p_2 x_2 + \ldots + p_N x_N\end{aligned}$$

- The expected value of a random variable is simply its mean

Expectations

- In the coin-flipping example, we can let the random variable x denote our winnings on a given turn
- x can be either 1 or -1: $X = \{-1, 1\}$
- The probability of either realization is specified by $\ell = (\frac{1}{2}, \frac{1}{2})$
- The average winnings, or expected value of x, is given by:

$$\mathbb{E}[x] = \frac{1}{2}(-1) + \frac{1}{2}(1) = 0$$

Expectations

- In summary, the *expected value* of a random variable x simply gives its mean, or average realization
- Often, utility depends directly on the realization of x
- For example, x may denote a random amount of dollars, and money impacts utility
- We can use expectations to compute the *expected utility* of a function $u(x)$
- Before doing this, we should define a couple of objects

Bernoulli Utility

- Given a set of events $X = \{x_1, \ldots, x_N\}$, the **Bernoulli** utility function $u(x)$ assigns a level of utility to each event:

$$u : X \to \mathbb{R}$$

- Works exactly like the “normal” utility functions we’ve seen before
- For example:
  - $X = \{sick, healthy\}$
  - $u(sick) = -10$, $u(healthy) = 10$
- $u(x)$ simply tells us our utility for any possible realization of x

Expected Utility

- Given a set of events X, the corresponding lottery ℓ, and a Bernoulli utility function $u(x)$ specifying utility from any event, we can compute our expected utility

- Expected utility is given by:

$$\begin{aligned} U(\ell) = \mathbb{E}[u(x)] &= \sum_{i=1}^{N} p_i u(x_i) \\ &= p_1 u(x_1) + p_2 u(x_2) + \ldots + p_N u(x_N) \end{aligned}$$

- $U(\ell)$ is called the Von-Neumann Morgenstern (VNM) utility function

- $u(x)$ assigns utility to events, while $U(\ell)$ assigns utility to lotteries

Taking Inventory

- Remember the distinction between *expected value* and *expected utility*
  - Expected value - $\mathbb{E}[x]$ is the mean (average) of x
  - Expected utility - $\mathbb{E}[u(x)]$ is the mean (average) of $u(x)$
- Remember the distinction between *Bernoulli* utility and *VNM* utility
  - Bernoulli - $u(x)$ assigns utility for each potential event
  - VNM - $U(\ell)$ assigns utility to each potential lottery
- When making decisions under uncertainty, we model agents as VNM utility maximizers
- How does risk/uncertainty impact optimal decision making? To answer this, let's first take a look at an important mathematical result.

Jensen's Inequality

Theorem

If $u(x)$ is strictly concave, then:

$$\mathbb{E}[u(x)] < u(\mathbb{E}[x])$$

If $u(x)$ is strictly convex, then:

$$\mathbb{E}[u(x)] > u(\mathbb{E}[x])$$

If $u(x)$ is linear, then:

$$\mathbb{E}[u(x)] = u(\mathbb{E}[x])$$

- Jensen's inequality relates functions with their expectations
- How is this related to risk preferences?

Risk Preferences

- Consider the following scenario: there is a box in front of you with an unknown amount of money
- The possible amounts of money are $X = \{1, 9\}$
  - i.e. the box has either \$1 or \$9
- Each value is equally probably (i.e. $\ell = (\frac{1}{2}, \frac{1}{2})$), so on average you'll draw \$5: $\mathbb{E}[x] = 5$
- You have two options:
  1. Open the box and collect whatever is inside
  2. Get \$5 and walk away
- Which option should you choose?

Risk Preferences

- If you take the money and walk away, you'll get $\mathbb{E}[x] = 5$ for sure
- Your utility will be $u(\mathbb{E}[x])$ with probability 1
- If you decide to open the box, your expected utility is $\mathbb{E}[u(x)]$
- Always, we do what gives the highest expected payoff
- Jensen's inequality states that whether $u(\mathbb{E}[x])$ or $\mathbb{E}[u(x)]$ is larger depends on the curvature of $u(x)$
- $u(x)$ can be one of three things:
  1. Linear
  2. Convex
  3. Concave

Linear Utility

- If $u(x)$ is linear ($u''(x) = 0$), then by Jensen's inequality:

$$u(\mathbb{E}[x]) = \mathbb{E}[u(x)]$$

- If $u(x)$ is linear, we say the agent is *risk-neutral*
- They are indifferent between taking the \$5 and opening the box
- Risk has no impact on their decision making

Convex Utility

- What if $u(x)$ is convex? i.e. $(u''(x) > 0)$
- Then by Jensen's inequality:

$$u(\mathbb{E}[x]) < \mathbb{E}[u(x)]$$

- If $u(x)$ is convex, we say the agent is *risk-loving*
- A risk-loving agent prefers opening the box to taking the \$5
- Risk-loving agents have preference for risky alternatives over safe ones

Concave Utility

- What about the last case where $u(x)$ is concave? i.e. $(u''(x) < 0)$
- By Jensen's inequality:

$$u(\mathbb{E}[x]) > \mathbb{E}[u(x)]$$

- If $u(x)$ is concave, we say the agent is *risk-averse*
- A risk-averse agent prefers taking the \$5 to opening the box
- Risk-averse agents have preference for safe alternatives over risky ones

Risk Preferences

- We've mentioned the three possible preferences over risk
  1. Risk-aversion
  2. Risk-lovingness
  3. Risk-neutrality
- Which category an agent falls into depends on the curvature of their Bernoulli utility function $u(x)$
- Let's look at some graphical representations of utility in each of the three cases mentioned above

Risk-Aversion

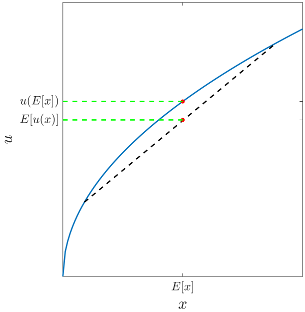

- If $u''(x) < 0$, then $\mathbb{E}[u(x)] < u(\mathbb{E}[x])$
- Notice for a risk-averse DM: $\frac{\partial MU_x}{\partial x} < 0 \rightarrow$ marginal utility diminishes

Risk-Aversion

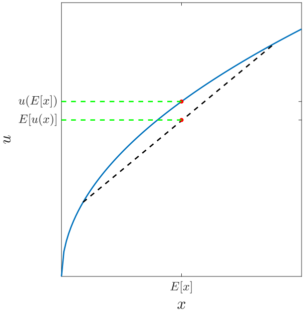

- Risk-averse agents hate losing more than they love winning
  - i.e. gaining \$1 increases utility by less than losing \$1 decreases it
- As a result, prefer certainty over uncertainty

Risk-Lovingness

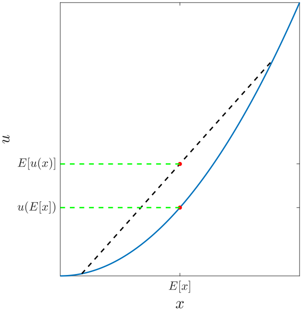

- If $u''(x) > 0$, then $\mathbb{E}[u(x)] > u(\mathbb{E}[x])$
- Notice for a risk-loving DM: $\frac{\partial MU_x}{\partial x} > 0 \rightarrow$ marginal utility increases

Risk-Lovingness

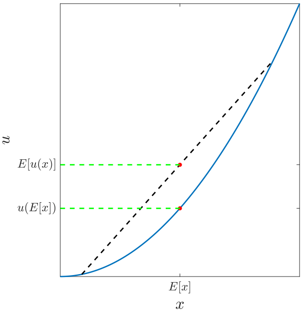

- Risk-loving agents love winning more than they hate losing
  - i.e. gaining \$1 increases utility by more than losing \$1 decreases it
- As a result, prefer uncertainty over certainty

Risk-Neutrality

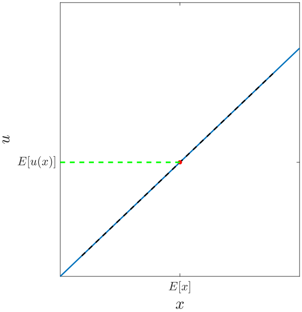

- If $u''(x) = 0$, then $\mathbb{E}[u(x)] = u(\mathbb{E}[x])$
- Notice for a risk-neutral DM: $\frac{\partial MU_x}{\partial x} = 0 \rightarrow$ constant marginal utility

Risk-Neutrality

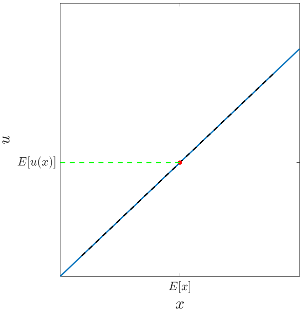

- Risk-neutral agents love winning the same amount that they hate losing
  - i.e. gaining \$1 increases utility by the same as losing \$1 decreases it
- As a result, indifferent between certainty and uncertainty

Taking Inventory

- Whether a person likes, dislikes, or is indifferent towards risk depends on the curvature of their utility
  - $u''(x) < 0 \rightarrow$ risk-averse
  - $u''(x) > 0 \rightarrow$ risk-loving
  - $u''(x) = 0 \rightarrow$ risk-neutral
- There is a close connection between risk preferences and marginal utility
- In each of these three cases, let's think about people's willingness to avoid/seek risk

Certainty Equivalent

- Back to our Pandora's box scenario
- We've determined that whether somebody prefers the $5 to opening the box depends on their risk preferences
  - i.e. depends on the curvature of their utility
- Let's now ask a different question: How much money could I offer you such that you would not open the box?
- In other words: how much for-sure money would it take such that you are indifferent between opening the box and walking away?

Certainty Equivalent

- Let me repeat the same question using math
- What is the monetary value, c, such that:

$$u(c) = \mathbb{E}[u(x)]$$

- The c which solves the equation above is called the *certainty equivalent*
- An agent is indifferent between getting c for sure (i.e. with probability 1) and "playing the lottery"

Certainty Equivalent

- A natural question: is c higher or lower than $\mathbb{E}[x]$?
- If $c < \mathbb{E}[x]$, the agent is willing to give up money to avoid risk
- If $c > \mathbb{E}[x]$, the agent must be paid to avoid risk
- Which case we fall into again depends on the curvature of $u(x)$
- Let's go through each of the three cases

Certainty Equivalent - Risk-Averse Agent

- Suppose an agent has utility function:

$$u(x) = \sqrt{x}$$

- $u''(x) < 0$, so this person is risk-averse

- The certainty equivalent c solves:

$$\begin{aligned}
u(c) &= \frac{1}{2}u(1) + \frac{1}{2}u(9) \\
\sqrt{c} &= \frac{1}{2}\sqrt{1} + \frac{1}{2}\sqrt{9} \\
\sqrt{c} &= \frac{1}{2}1 + \frac{1}{2}3 \\
c &= 4
\end{aligned}$$

Certainty Equivalent - Risk-Averse Agent

- Given the utility function $u(x) = \sqrt{x}$, we determined that the certainty equivalent of the Pandora's box game is $c = 4$
- Here, the certainty equivalent is less than the expected value
  - Recall that $\mathbb{E}[x] = 5$
- In fact, $\mathbb{E}[x] > c$ always holds for risk-averse agents
- Risk-averse agents are willing to sacrifice money (in expectation) in order to avoid risk

Certainty Equivalent - Risk-Loving Agent

- Suppose an agent has utility function:

$$u(x) = x^2$$

- $u''(x) > 0$, so this person is risk-loving
- The certainty equivalent c solves:

$$\begin{aligned} u(c) &= \frac{1}{2}u(1) + \frac{1}{2}u(9) \\ c^2 &= \frac{1}{2}1^2 + \frac{1}{2}9^2 \\ c^2 &= \frac{1}{2}1 + \frac{1}{2}81 \\ c &= \sqrt{41} \approx 6.4 \end{aligned}$$

Certainty Equivalent - Risk-Loving Agent

- Given the utility function $u(x) = x^2$, we determined that the certainty equivalent of the Pandora's box game is $c \approx 6.4$
- Here, the certainty equivalent is higher than than the expected value
- For risk-loving agents, it is always the case that $\mathbb{E}[x] < c$
- Risk-loving agents must be paid to in order to avoid taking risks

Certainty Equivalent - Risk-Neutral Agent

- Suppose an agent has utility function:

$$u(x) = x$$

- $u''(x) = 0$, so this person is risk-neutral

- The certainty equivalent c solves:

$$u(c) = \frac{1}{2}u(0) + \frac{1}{2}u(1)$$
$$c = \frac{1}{2}1 + \frac{1}{2}9$$
$$c = 5$$

Certainty Equivalent - Risk-Neutral Agent

- Given the utility function $u(x) = x$, we determined that its certainty equivalent is $c = 5$
- Here, the certainty equivalent is equal to the expected value
- For risk-neutral agents, it is always the case that $\mathbb{E}[x] = c$
- Risk-neutral agents must be paid the expected value of a lottery in exchange for not playing it
- They are indifferent between getting \$5 for sure or taking a gamble which yields \$5 on average

Certainty Equivalent - Summary

- In summary:
  - Risk-averse: $c < \mathbb{E}[x]$
  - Risk-loving: $c > \mathbb{E}[x]$
  - Risk-neutral: $c = \mathbb{E}[x]$
- If an agent is risk-averse:
  - Willing to sacrifice money in exchange for avoiding risk
- If an agent is risk-loving:
  - Must be paid money in exchange for avoiding risk
- If an agent is risk-averse:
  - Indifferent between risk and no risk

Risk Premium

- Using our Pandora's box example, let's ask a slightly different question
- How much money would you be willing to sacrifice to avoid playing the game?
- We can answer this question by comparing the expected value of the game, $\mathbb{E}[x]$, with the agent's certainty equivalent c
- The difference between $\mathbb{E}[x]$ and c is called the *risk premium* (r):

$$r = \mathbb{E}[x] - c$$

Risk Premium - Risk-Averse Agent

- If an agent is risk-averse ($u''(x) < 0$), then we known that:

$$\mathbb{E}[x] > c$$

- As a result, risk-averse agents have a positive risk premium:

$$r = \mathbb{E}[x] - c > 0$$

- Again, willing to sacrifice money to avoid playing the lottery

Risk Premium - Risk-Loving Agent

- If an agent is risk-loving ($u''(x) > 0$), then we known that:

$$\mathbb{E}[x] < c$$

- As a result, risk-loving agents have a negative risk premium:

$$r = \mathbb{E}[x] - c < 0$$

- Must be paid to avoid playing a lottery

Risk Premium - Risk-Neutral Agent

- If an agent is risk-averse ($u''(x) = 0$), then we known that:

$$\mathbb{E}[x] = c$$

- As a result, risk-averse agents have a risk premium equal to zero:

$$r = \mathbb{E}[x] - c = 0$$

- Indifferent between lottery & no lottery, no need to pay them

Computing the Risk Premium

- How do we actually compute the risk premium?
- Recall the formula for the certainty equivalent:

$$u(c) = \mathbb{E}[u(x)]$$

- Since $r = \mathbb{E}[x] - c$, then:

$$u(\mathbb{E}[x] - r) = \mathbb{E}[u(x)]$$

- The risk premium r can be computed using the formula above
- Let's work through an example

Computing the Risk Premium

- Suppose we're given the utility function:

$$u(x) = x^{1/2}$$

- With .25 probability, we win \$16. With .75 probability, we win \$0.
  - $X = \{0, 16\}$
  - $\ell = (.75, .25)$

- What is the risk premium in this case?

- We simply apply the formula:

$$u(\mathbb{E}[x] - r) = \mathbb{E}[u(x)]$$

Computing the Risk Premium

$$u(\mathbb{E}[x] - r) = \mathbb{E}[u(x)]$$

- We need to compute $\mathbb{E}[x]$, $\mathbb{E}[u(x)]$, then plug everything into the formula
- The expected value in this case:

$$\mathbb{E}[x] = .75(0) + .25(16) = 4$$

- The expected utility in this case:

$$\begin{aligned}\mathbb{E}[u(x)] &= .75u(0) + .25u(16) \\ &= .75(0)^{1/2} + .25(16)^{1/2} = 1\end{aligned}$$

Computing the Risk Premium

$$u(\mathbb{E}[x] - r) = \mathbb{E}[u(x)]$$

- Plugging $\mathbb{E}[x] = 4$ and $\mathbb{E}[u(x)] = 1$ into the formula:

$$\begin{aligned} u(4 - r) &= 1 \\ (4 - r)^{1/2} &= 1 \\ 4 - r &= 1 \\ r &= 3 \end{aligned}$$

- A person with utility function $u(x) = x^{1/2}$ would accept \$3 less than the expected value of this lottery in order to avoid playing it
- Note that we alternatively could have just computed the certainty equivalent then applied the formula: $c = \mathbb{E}[x] - r$

Measuring Risk-Aversion

- For many reasons, we most often to think about people as being risk-averse
- It would be useful to have a measure of "how risk-averse" a person is
- The Arrow-Pratt coefficient of absolute risk aversion $A(x)$ does this for us:
$$A(x) = -\frac{u''(x)}{u'(x)}$$
- $A(x) \in (-\infty, \infty)$, meaning it can take any real positive or negative value
- The higher $A(x)$, the more risk-averse a person is
- The lower $A(x)$, the less risk-averse they are

Measuring Risk-Aversion

- In addition, it is possible that a person's level of risk aversion changes with their level of wealth
  - i.e. as they become more wealthy, they become more/less risk-averse
- If $A'(x) > 0$, the agent exhibits *increasing absolute risk aversion*
  - The wealthier they are, the more risk-averse they are
- If $A'(x) < 0$, the agent exhibits *decreasing absolute risk aversion*
  - The wealthier they are, the less risk-averse they are
- If $A'(x) = 0$, the agent exhibits *constant absolute risk aversion*
  - As they become wealthier, they maintain the same level of risk-aversion

Risk & Insurance

- If a person is risk-averse (i.e. $u''(x) < 0$), we've demonstrated that they'd be willing to pay money to avoid risk
- Naturally, we see markets arise which offer people protection against risk (for a price)
- In particular, firms offer *insurance* to people in exchange for protection against risk
- Insurance pops up in many forms:
  - Health insurance
  - Car insurance
  - Home insurance
  - Life insurance
  - Deposit insurance
  - Etc...

Risk & Insurance

- How should firms price their insurance plans?
- Part of answering this question is first determining how much consumers are willing to pay for insurance
- A person's willingness to pay for insurance is called their *insurance premium*, and is denoted by i
- It is very similar to the risk premium
- Let's go through a simple example of an insurance premium computation

Insurance Premium

- Suppose that a person has $W = \$144$ in wealth, and has the utility function:
$$u(x) = x^{1/2}$$
- They face the following scenario:
  - With probability .25, they get sick and incur \$44 in medical expenses
  - With probability .75, they don't get sick and incur \$0 in medical expenses
- If they get sick, their final wealth is $144 - 44 = 100$
- If they don't get sick, their final wealth remains at \$144

Insurance Premium

- The set of possible events is: $X = \{100, 144\}$
  - i.e. they end up with either \$100 or \$144
- The corresponding lottery is $\ell = (.25, .75)$
- What is their willingness to pay for *perfect* insurance?
  - Perfect insurance: absolutely no risk upon buying it
- The highest amount they'd be willing to pay for insurance, their insurance premium (i), makes them indifferent between insurance and no insurance:

$$u(W - i) = \mathbb{E}[u(x)]$$

- We can use this formula to compute the person's insurance premium

Insurance Premium

$$u(W - i) = \mathbb{E}[u(x)]$$

- The expected utility in this case is:

$$\begin{aligned}\mathbb{E}[u(x)] &= .25(100)^{1/2} + .75(144)^{1/2} \\ &= .25(10) + .75(12) \\ &= 2.5 + 9 = 11.5\end{aligned}$$

- Plugging this into the insurance premium formula:

$$\begin{aligned}u(W - i) &= \mathbb{E}[u(x)] \\ (144 - i)^{1/2} &= 11.5 \\ 144 - i &= 132.25 \\ i &= 11.75\end{aligned}$$

Insurance Premium

- In this example, the person is willing to pay at most $i = 11.75$ for perfect insurance
- The insurance premium came out to be positive because this person was risk-averse
- If a person is risk-loving: negative insurance premium
- If a person is risk-neutral: insurance premium is zero
- But is \$11.75 a "fair" price for insurance here?
- Let's talk about what fairness means in the context of insurance

Insurance Markets

- In general, an insurers expected profit is given by:

$$\mathbb{E}[\pi] = y - pq$$

- y is the insurance premium they charge (which they get for sure)
- q is the amount they pay in case of an adverse event which occurs with probability p
  - Ex: amount insurer pays following a medical event, car accident, etc.
- In the example we're working with, $p = .25$ and $q = 44$, so:

$$\mathbb{E}[\pi] = y - .25(44)$$

$$\mathbb{E}[\pi] = y - .25(44)$$

- If the insurer charges the consumer's insurance premium $y = i = 11.75$, then their expected profits are:

$$\mathbb{E}[\pi] = 11.75 - .25(44) = 11.75 - 11 = .75 > 0$$

- The insurer is profiting at the expense of the consumer
    - Could have lowered the cost of insurance and made the consumer better off
- The rate $y = 11.75$ is *actuarially unfair*

Actuarial Fairness

- An insurance rate is *actuarially fair* if it yields the insurer zero expected profits
- In the previous example, the actuarially fair rate is:

$$\begin{aligned} \mathbb{E}[\pi] &= 0 \\ y - .25(44) &= 0 \\ y^* &= 11 \end{aligned}$$

- In general, the actuarially fair rate is:

$$\begin{aligned} \mathbb{E}[\pi] &= 0 \\ y - pq &= 0 \\ y^* &= pq \end{aligned}$$

- Actuarially fair insurance rates exactly offset expected payouts

Actuarial Fairness

- A few observations are in order
- i, the consumer's insurance premium, is the max they're WTP for insurance
- y^*, the actuarially fair rate, is the min a firm would charge for insurance
- If $i > y^*$: firm is willing to offer insurance
- If $i < y^*$: firm is not willing to offer insurance
  - Risk-averse consumer desires insurance, but their WTP for insurance is too small for the insurer to make any profits

Screening & Signaling
or
“Hidden Type”

Screening & Signaling

- In many settings, agents may have some *private information*
- For example:
  - Consumers have better knowledge of their health than insurers do
  - Workers have better knowledge of their ability than employers do
  - Salesmen have better knowledge of their product's quality than consumers do
- Even if I am uncertain about somebody's characteristics (i.e. their "type"), I can learn a lot about who they are by observing their decisions
  - Their choices send a *signal* of their type
- Let's go through an example of this sort of scenario

Screening & Signaling

- Why is getting a college degree valuable?
- One possibility: going to college allows you to build skills which are valuable in the labor market
  - i.e. college has a direct (and positive) impact on your productivity
- Another possibility: going to college signals to employers that you are of high ability
  - i.e. people who go to college are inherently skilled, college has no direct impact on productivity
- The signaling hypothesis was popularized by **Michael Spence**
- Which of the two hypotheses is true? In reality, its probably a mixture of both

Screening & Signaling

- Let's suppose for simplicity that college only has signaling value
- Let's consider a worker's decision over whether to go to college
- The worker is either a "high type" $(t = H)$ or a "low type" $(t = L)$
- The labor market does not know the type of the worker
- However, the worker can go to college to try to signal their type to potential employers
  - Employers assign higher probability to $t = H$ if the worker has a degree
- Assume that expected wages with a degree are higher than expected wages with no degree:

$$\mathbb{E}[\, w \mid college \,] > \mathbb{E}[\, w \mid no\ college \,]$$

Screening & Signaling

- Let's assume that going to college costs c
  - c can reflect both tuition costs and "non-pecuniary costs" (ex. stress, anxiety, etc.)
- When is it worthwhile to go to college?
- The worker goes to college if:

$$\mathbb{E}[\, w \mid \textit{college} \,] - c > \mathbb{E}[\, w \mid \textit{no college} \,]$$
$$\mathbb{E}[\, w \mid \textit{college} \,] - \mathbb{E}[\, w \mid \textit{no college} \,] > c$$

- It is optimal to go to college if the wage premium covers the cost of attendance
- However, there is a problem

Screening & Signaling

- If both the high types and low types face the same cost of obtaining a degree, either both types attend or both types don't
- In this case, college has no signaling value
  - There is a *pooling equilibrium*
- For college to have signaling value, costs must be such that high types attend, but low types don't
- Let's assume now that the cost of college attendance is:
  - c_H if $t = H$
  - c_L if $t = L$
  - $c_H < c_L$
- Now, it is less costly for high types to get degrees than low types

- High types get a degree if:

$$\mathbb{E}[\, w \mid college \,] - \mathbb{E}[\, w \mid no\ college \,] > c_H$$

- To prevent low types from getting degrees, we need:

$$\mathbb{E}[\, w \mid college \,] - \mathbb{E}[\, w \mid no\ college \,] < c_L$$

- As long as the wage premium lies in the range (c_L, c_H), high types go to college while low types don't
    - We have a *separating equilibrium*
- Wage premium too high $\rightarrow$ everyone goes to college
    - College reveals nothing about the worker's type
- Wage premium too low $\rightarrow$ nobody goes to college
    - Again, no information is transmitted

Screening & Signaling

- For people's decisions to convey any information about them, we need some notion of *separation* to hold
- In other words, we need people of type H to choose a, and people of type L to choose b
- Then, I can infer your type based upon your behavior
- If all types make the same decisions, observing your behavior reveals nothing

Moral Hazard
or
“Hidden Action”

Moral Hazard

- Alternatively, I may have perfect information about somebody's type, but I may be uncertain about their "action"
- I don't know their action (i.e. decision), but their decision is relevant for my payoff
- For example:
  - Managers don't observe worker effort, which impact manager profits
  - Medical insurers don't observe insurees' health behavior, which impacts medical expenditures
  - FDIC may not observe banks' decisions over portfolio risk, which impacts expected deposit insurance payouts
- Let's talk through an example of this

Moral Hazard

- Let's consider two parties:
  - The *principal* (the "manager")
  - The *agent* (the "worker")
- The agent makes decisions which directly impacts the principal's payoff
  - Ex: worker chooses how much effort to exert, which impacts the manager's profits
- The principal would like the agent to select a particular action
  - Ex: a manager wants their worker to exert as much effort as possible
- But, the principal does not directly observe the agent's decisions
  - The agent has some private information

Moral Hazard

- Let's assume the manager's profits y are given by:

$$y = e + \epsilon$$

- $e \in \{0, 1\}$ denotes the worker's effort choice
  - $e = 1$ denotes "high effort"
  - $e = 0$ denotes "low effort"

- $\epsilon \in \{0, 1\}$ is a random shock
  - $\epsilon = 1$ with probability $\frac{1}{2}$
  - $\epsilon = 0$ with probability $\frac{1}{2}$

- The manager's expected profit is:

$$\mathbb{E}[y] = e + \frac{1}{2}$$

Moral Hazard

$$\mathbb{E}[y] = e + \frac{1}{2}$$

- Manager profits increase in worker effort, so they'd like effort to be as high as possible
- But, they don't directly observe effort. How can they induce the worker to work hard?
- They can design a compensation package x which incentivizes high effort

Moral Hazard

- Suppose the worker's utility is given by:

$$u(x, e) = x^{1/2} - ce$$

- The worker is risk-averse, so prefers certain pay to uncertain pay
- Additionally, they dislike exerting effort
  - If they exert high effort ($e = 1$), then they incur cost $c > 0$
- Summary so far:
  - Manager does not observe worker effort
  - Effort increases profits, so manager wants high effort
  - Effort is costly for the worker, so they would prefer low effort over high effort
  - Despite this, manager can financially incentivize the worker to exert high effort

Moral Hazard

- How should compensation be designed?
- First thought: worker is risk-averse, prefers certain pay over uncertain pay, so let's pay them a fixed wage $x = w$
- In this case, worker's expected payoff from exerting high effort is:

$$\begin{aligned}\mathbb{E}[\, u(x,e) \mid e = 1\,] &= \frac{1}{2}\sqrt{w} + \frac{1}{2}\sqrt{w} - c \\ &= \sqrt{w} - c\end{aligned}$$

- Expected payoff from exerting low effort is:

$$\mathbb{E}[\, u(x,e) \mid e = 0\,] = \frac{1}{2}\sqrt{w} + \frac{1}{2}\sqrt{w} = \sqrt{w}$$

Moral Hazard

$$\mathbb{E}[\,u(x,e)\,|\,e=0\,] > \mathbb{E}[\,u(x,e)\,|\,e=1\,]$$

- With a fixed wage, the worker is always better off choosing low effort over high effort
- Fixed wage does not incentivize the worker
- Consider an alternative compensation package which directly ties pay to firm profits:
$$x = \beta y$$
- β is the *pay-performance sensitivity*

Moral Hazard

- If $x = \beta y$, when does the worker exert high effort?
- High effort is optimal if:

$$\begin{aligned} \mathbb{E}[\, u(x,e) \mid e = 1\,] &> \mathbb{E}[\, u(x,e) \mid e = 0\,] \\ \frac{1}{2}\sqrt{2\beta} + \frac{1}{2}\sqrt{\beta} - c &> \frac{1}{2}\sqrt{\beta} + \frac{1}{2}\sqrt{0} \\ \frac{1}{2}\sqrt{2\beta} &> c \\ \beta &> 2c^2 \end{aligned}$$

- If $\beta > 2c^2$, then high effort is optimal for the worker
- If financial incentives (i.e. pay sensitivity) are high enough, then high effort can be induced

Moral Hazard

- Fixed wage offered no incentives for the worker
- But if pay is sufficiently sensitive to performance, it is in the worker's interest to exert high effort
- Remember the worker is risk averse. If their pay is tied to y, which is random, the worker wants to do all they can to ensure y is as high as possible.
  - They hate "losing," so want to prevent this at all costs
- Properly designed compensation packages can help align incentives between firms and their workers
- More generally: if you want to induce somebody to do a over b, increase their payoff from a and decrease their payoff from b