

Decision Making under Uncertainty

by Kevin Hasker

5, February, 2019

1 An Introduction, and some simple facts.

First of all, since we are looking at a complex phenomena in many dimension we are going to simplify it in one by only considering the outcome of money. Everything will be basically boiled down into cash. Cash lost, cash won. And before we get into further analysis, we need to get more specific about what we are considering. Thus:

Definition 1 A lottery, (L) is a set of outcomes and probabilities of such outcomes, $L = (p, x) = (p_i, x_i)_{i=1}^n$.

For example, say that with probability $\frac{1}{2}$ you win 10,000 TL (10^4) and with probability $\frac{1}{2}$ you loose 10,000 TL, we would formally write this as $L_f = ((\frac{1}{2}, 10,000), (\frac{1}{2}, -10,000))$. Another simple lottery would be I give you nothing for sure, $L_0 = (1, 0)$. A more complex lottery might be I give you 2^n TL with probability $1/2^n$, we would write this as: $L_{sp} = (\frac{1}{2^n}, 2^n)_{n=1}^\infty$. The most general lotteries will be defined over the reals, so the outcome space is always just the reals, and these can be described purely by their cumulative probability function (*CPF*) denoted F . Formally speaking: $F(x) = \Pr(\text{the outcome is } x \text{ or lower})$.

A word on terminology, some people for religious or other reasons have an attitude that gambling, and therefore lotteries, are bad. Please don't fall into the trap of thinking that this is about gambling. It is nothing more than the real risk we face. For example, what is the probability you will have a car accident in the next year? You probably, with some work, could figure this out. That makes it a lottery. You are not certain to have a car accident, rather you might have one so we would call that a lottery. As well, what is the inflation rate for next year going to be? I don't really care about money, so I don't know, but someone who does care could probably tell you what it is most likely to be, and what is the probability it is (say) 2% higher or lower than that. Again, we would refer to this as a lottery.

The simplest utility function in this setting would clearly be just the expectation, expected value, or the average amount. In other words for a lottery $L = (p, x) = (p_i, x_i)_{i=1}^n$ it would be $U(L) = \sum_{i=1}^n p_i x_i$. To give a concrete example,

$$U(L_f) = \frac{1}{2}(10,000) + \frac{1}{2}(-10,000) = 0 = U(L_0) .$$

But a rather clever example invented by Bernoulli shows this simple utility function will not work.

Example 2 (Saint Petersburg Paradox) *I will flip a fair coin. If it comes up heads, I give you 2 TL, if it comes up tails I flip again. If it comes up head*

on the second flip, I give you 4 TL, if it comes up tails I flip again. In general, if it comes up heads on the n 'th flip I give you 2^n TL, if it comes up tails I flip again.

How much would you pay for this lottery? 2 TL seems like a sensible minimum, because that's the least you could win. Would you pay 100 TL? How about 200TL? Anyone ready to pay 1000TL?

Let me point out how high the winnings might get. If you are lucky and the coin doesn't hit heads until the 10'th round you would win $2^{10} = 1024$ TL. If you get to 20 then you will get $2^{20} = 1,048,576$ TL—over a million!

The funny thing is that, by construction, someone who's willing to pay the expected value will give an infinite amount of cash to play this.

In order for you to win in round n , you have to have had a tail every previous period which happens with probability $1/2^{n-1}$, then you have to have a head in round n , which happens with probability $1/2$, and you will win 2^n lira. Thus:

$$E(n) = \left(\frac{1}{2^{n-1}}\right) \left(\frac{1}{2}\right) 2^n = 1 .$$

Now this is true for every n , so

$$E(L_{sp}) = \sum_{n=1}^{\infty} E(n) = \sum_{n=1}^{\infty} 1 \rightarrow \infty$$

pretty clever, huh?

Based on this Bernoulli proposed that people evaluate lotteries using a *Bernoulli utility function* $u(m)$ where m is the amount of money. It is quite simple to show that $U(L_{sp}) = \sum_{n=1}^{\infty} \frac{1}{2^n} u(2^n) < \infty$ can easily be satisfied, or that people would be willing to pay a finite amount for this lottery.

2 The Theory of Expected Utility

Notice that Bernoulli was basically proposing that utility would be the weighted sum of the utility of the outcomes, a model we now call expected utility. This model, however, was not formalized until the second edition of *Von Neumann and Morgenstern (1947) "The Theory of Games and Economic Behavior."* The primary topic of this book was Game Theory, or a theory of interactions, but in order to understand behavior they needed to understand unpredictable behavior. There are so many interactions (penalty shots in soccer, serving in tennis, the game rock/paper/scissors) where it is obviously optimal to be unpredictable that we could no longer ignore the fact that people and the world might not be predictable. Their model, formally known as *Von Neumann-Morgenstern Utility* (VnM utility) was what we wrote above:

$$U(L) = \sum_{i=1}^n p_i u(x_i) .$$

Notice it is based on a *Bernoulli* utility function $u(\cdot)$ which is defined over the outcome, it is different because it is defined over lotteries. To get this representation they had to make one strong, key assumption:

Definition 3 *The independence axiom says that if $x \sim y$ (I like x as much as y) then for any $p \in [0, 1]$ the lottery $L = ((p, x), (1 - p, y))$ is as good as y , or $L \sim y$.*


This implies that for any z , the lottery $((qp, x), (q(1 - p), y), (1 - q, z))$ is as good as the lottery $((q, y), (1 - q, z))$. And here I can give a pretty simple counter-example. Say you are indifferent over a trip to China or a trip to India. Then I agree that you wouldn't mind flipping a coin to decide which to go to. Formally, if $x = \{\text{trip to India}\}$ $y = \{\text{trip to China}\}$ then I accept that $((\frac{1}{2}, x), (\frac{1}{2}, y)) \sim y$. Now consider a third outcome, $z = \{\text{Watching a film about a trip to India}\}$. Do you now think that $((\frac{1}{3}, x), (\frac{1}{3}, y), (\frac{1}{3}, z)) \sim ((\frac{2}{3}, y), (\frac{1}{3}, z))$? Me personally I would say that the second lottery—going to China 2/3 of the time and watching a film about a trip to India one third of the time—is strictly better than a lottery where I might lose my chance to go to India and then have to watch a movie about going to India. That would stink. I would hate that.

But while this "axiom" is anything but self evident (as axioms are supposed to be) and we can show clear violations of it in many interesting settings it is still a good basic assumption. As someone familiar with the research on violations of this axiom, this research is interesting but generally only matters on the fringes of our analysis. This simple model is our workhorse, and will continue to be for the foreseeable future.

2.1 Risk Aversion

Risk Aversion merely means that you would rather get the average or expected outcome with certainty than a lottery. For example, would you be willing to take the lottery $L_f = ((\frac{1}{2}, 10,000), (\frac{1}{2}, -10,000))$ instead of the lottery $L_0 = (1, 0)$? It's fine if you would be, but most people would not. Possibly losing 10,000 TL over the flip of a coin? That... well that seems very risky. Yea, sure, you might win 10,000 TL but still. Now let us plot the utility of these two outcomes. We want that utility is increasing in wealth, so $u(W + 10^4) > u(W - 10^4)$ where W is your current wealth in TL. Other than that, the points are completely arbitrary. In the graph I have $3 = W - 10^4$ and $7 = W + 10^4$ so we know that $W = 5$. Now if you agree that $U(L_0) > U(L_f)$ your utility can not be in the dotted region, because $\frac{1}{2}(W - 10^4) + \frac{1}{2}(W + 10^4) = W$ but you have $\frac{1}{2}u(W - 10^4) + \frac{1}{2}u(W + 10^4) < u(W)$. Indeed your $u(W)$ must be in the solid

line area if you think $U(L_0) > U(L_f)$.



What type of function has the characteristic that $U(L_0) > U(L_f)$? It's quite simple, a concave one.

Definition 4 A function is concave if for any α , $0 \leq \alpha \leq 1$ $u(\alpha x + (1 - \alpha)y) \geq \alpha u(x) + (1 - \alpha)u(y)$ it is strictly concave if the inequality is strict when $0 < \alpha < 1$.

Something useful to recognize is the essential equivalence of risk aversion and concavity. For a general lottery F , risk aversion means:

$$u\left(\int_x x dF(x)\right) \geq \int_x u(x) dF(x)$$

this is known as *Jensen's inequality*.

To be explicit we will always assume:

Assumption If $u(W)$ is our Bernoulli utility function, then $u'(W) > 0$ and $u''(W) < 0$ if the person is strictly risk averse, $u''(W) = 0$ if and only if the person is risk neutral, and otherwise the person is strictly risk loving.

Just to be clear, unfortunately people are not always risk averse or risk loving. Almost everyone has or wants some form of insurance or other, and yet many people enjoy going to the casino and gambling. I don't, but that's because I don't like money that much. Even more people buy lottery tickets. In case you didn't know, in both of these cases people are going to loose money on

average. With casinos, well obviously the house has to pay its costs of running the casino and wants to make some money on top. With lotteries, in the United States these are almost all run by governments and they brag about how the money they earn will build schools and other things (it never happens, but still.) The point is that in both cases they are making money. Therefore you must be losing money on average. Don't be fooled.

So sometimes people are both risk loving and risk averse. That is, of course, a problem for us and our simple model but outside the range of discussion in this class. We will assume people are risk averse.

2.2 Characterizing Preferences over Risk

We want to be able to compare people, say that person x is more risk averse than person y . To do this we must characterize risk. There are several different ways to do this.

2.2.1 Certainty Equivalence

The most important in the real world is the *certainty equivalence*.

Definition 5 *The certainty equivalence at the wealth level W of a lottery F is the function $c(W, F)$ such that:*

$$u(c(W, F)) = \int_x u(W + x) dF(x)$$

if a person is risk averse then $c(W, F) \leq \int_x (W + x) dF(x)$

The importance of this measure that $c(W, F) = W - P(W, F)$ where $P(W, F)$ is the insurance premium this person will pay to avoid this risk. This is something insurance companies have to calculate every day, and thus is vitally important to them.

While this measure is fundamentally and obviously important, the problem with this measure is that it obviously depends on what kind of risk we are facing. For different types of lotteries we do have a one dimensional characterization of the difference between these lotteries, but how will this measure change when we change lotteries? We would have to look at each one in particular.

2.2.2 Probability Premium

Perhaps we can get further if we focus on one particular type of lottery. We know that every risk averse person will reject the lottery $((\frac{1}{2}, x), (\frac{1}{2}, -x))$ so perhaps we should consider how much we have to shift the odds in their favor, or consider the lottery $L_\pi = ((\frac{1}{2} + \pi, x), (\frac{1}{2} - \pi, -x))$ and find the π such that:

$$u(W) = \left(\frac{1}{2} + \pi\right) u(W + x) + \left(\frac{1}{2} - \pi\right) u(W - x)$$

which has the solution

$$\pi(W, x) = \frac{\left[\frac{1}{2}u(W+x) + \frac{1}{2}u(W-x)\right] - u(W)}{u(W+x) - u(W-x)}.$$

This is called the *probability premium*.

Definition 6 *The probability premium is the amount you have to shift odds in your favor so that you are indifferent between $L_\pi = \left(\left(\frac{1}{2} + \pi, x\right), \left(\frac{1}{2} - \pi, -x\right)\right)$ and $L_0 = (1, 0)$.*

This is a second good measure, but there the problem is that it still depends on x . Furthermore while our first measure has a real world application it seems hard to imagine a real world application of this measure. So, the search continues.

2.2.3 Absolute Risk Aversion

The Arrow-Pratt coefficient of absolute risk aversion, or absolute risk aversion for short, is not immediately derived from any lottery. It has the benefit of being very simple.

Definition 7 *Absolute risk aversion is:*

$$r(W) = -\frac{u''(W)}{u'(W)}.$$

If a person is risk averse we know that $u''(W) \leq 0$ so $r(W) \geq 0$. How did we derive this? Well, consider a lottery $\sigma\varepsilon$ where $E(\varepsilon) = 0$ and $E(\varepsilon^2) = 1$, then we want to consider very small σ and find the (normalized) insurance premium of this lottery or:

$$u\left(W - P\frac{\sigma^2}{2}\right) = \int u(W + \varepsilon) d\varepsilon.$$

How do we solve for P ? We use the Taylor expansion.

Definition 8 *Assume that f has an infinite number of derivatives, then for any x and y :*

$$f(y) = f(x) + f'(x)(y-x) + f''(x)\frac{(y-x)^2}{2} + f'''(x)\frac{(y-x)^3}{3!}$$

or if $f^{(n)}(x)$ is the n 'th derivative of $f(\cdot)$ at x :

$$f(y) = \sum_{n=0}^{\infty} f^{(n)}(x) \frac{(y-x)^n}{n!}.$$

This is one of the most powerful results in mathematics for the empiricist. The key thing is to recognize that if y is very close to x then $\frac{(y-x)^n}{n!}$ is very close to zero. Thus close to x the function is essentially linear, a little further it is quadratic, and so on.

Here we can easily see that:

$$\begin{aligned} u\left(W - P\frac{\sigma^2}{2}\right) &= u(W) + \left(\left(W - P\frac{\sigma^2}{2}\right) - W\right) u'(W) + \tau(\sigma) \\ &= u(W) - P\frac{\sigma^2}{2} u'(W) + \tau(\sigma) \end{aligned}$$

the first term in $\tau(\sigma)$ is $u''(W) \left(\frac{P\sigma^2}{2}\right)^2 = u''(W) \frac{1}{16} P^2 \sigma^4$ so for small σ this is very small indeed. On the other hand we have

$$\begin{aligned} E[u(W + \varepsilon)] &= E\left[u(W) + \sigma\varepsilon u'(W) + \frac{\sigma^2\varepsilon^2}{2} u''(W) + \delta(\sigma)\right] \\ &= u(W) + \sigma E[\varepsilon] u'(W) + E[\varepsilon^2] \frac{\sigma^2}{2} u''(W) + E[\delta(\sigma)] \end{aligned}$$

and again, the first term in $E[\delta(\sigma)]$ is $E[\varepsilon^3] \frac{\sigma^3}{6} u'''(W)$, so for small σ it is very small again. Given that we assumed $E[\varepsilon] = 0$ and $E[\varepsilon^2] = 1$ this means:

$$E[u(W + \varepsilon)] = u(W) + \frac{\sigma^2}{2} u''(W) + E[\delta(\sigma)] .$$

so our original expression now is:

$$\begin{aligned} u\left(W - P\frac{\sigma^2}{2}\right) &= E[u(W + \varepsilon)] \\ u(W) - P\frac{\sigma^2}{2} u'(W) + \tau(\sigma) &= u(W) + \frac{\sigma^2}{2} u''(W) + E[\delta(\sigma)] \\ -P\frac{\sigma^2}{2} u'(W) &= \frac{\sigma^2}{2} u''(W) + E[\delta(\sigma)] - \tau(\sigma) \\ P\frac{\sigma^2}{2} u'(W) &= -\frac{\sigma^2}{2} u''(W) + \tau(\sigma) - E[\delta(\sigma)] \\ P &= -\frac{u''(W)}{u'(W)} + \frac{2}{\sigma^2} \frac{\tau(\sigma) - E[\delta(\sigma)]}{u'(W)} \end{aligned}$$

and as σ gets very small the last term: $2(\tau(\sigma) - E[\delta(\sigma)]) / (\sigma^2 u'(W))$ converges to zero, thus we get:

$$r(W) = -\frac{u''(W)}{u'(W)} = P$$

2.3 The Equivalence of all Three Measures and "More Risk Averse Than."

So... we have now constructed three lovely measures of risk aversion. The first is useful, the second less useful but simpler, the third even (potentially) least useful and hardly derived from an interesting lottery at all. So... what was the point? Oh yea, how do these three measures relate to "more risk averse than."

Definition 9 *Person A is more risk averse than person B if anytime A will choose a lottery over a sure outcome so will B.*

There might be risks B will face that A will not, but the reverse will never be true.

One of our most stunning results in this field is that all three of these measures are equivalent.

Theorem 10 *A is more risk averse than B if and only if:*

1. For all W and F , $c_A(W, F) \leq c_B(W, F)$
2. For all W and x , $\pi_A(W, x) \geq \pi_B(W, x)$
3. For all W $r_A(W) \geq r_B(W)$

Of course I will never ask you to prove this, but this tells us all of these measures are equivalent. There is no uncertainty about the best way to describe uncertainty, all the measures are the same. Of course this means we want to use the most parsimonious one, absolute risk aversion. Given our interest in this let us explore it a bit more.

2.4 The Relationship between Risk Aversion and Wealth, and Relative Risk Aversion.

First of all, let us find the coefficient of absolute risk aversion for several utility functions. First of all, the most common utility function used in finance is the quadratic utility function:

$$u(W) = \alpha W - \beta W^2 ,$$

where $\alpha > 0$ and $\beta > 0$. This utility function is so popular because $E(u(W)) = \alpha E(W) - \beta E(W^2)$, so these people care only about the expectation and variance—two very simple statistics. However it is ridiculous because:

$$r(W) = \frac{2\beta}{\alpha - 2\beta W} .$$

which is *increasing in wealth*. In other words with these preferences when someone gets richer they get more risk averse. Does this make sense? If you have a million TL a month a 1000 TL bet is more dangerous than someone who only has 500 TL spending money this month? I don't think so.

Stylized Fact If $W_A < W_B$ then person A should be more risk averse than person B .

So don't use this utility function, unless you're feeling lazy.

Another very interesting utility function is the *Constant Elasticities of Substitution* utility function. Like always this is favorite, when we have a one good utility function this is written as:

$$u(W) = \begin{cases} \frac{W^{1-\sigma}}{1-\sigma} & \sigma \neq 1 \\ \ln W & \sigma = 1 \end{cases}$$

notice that $\sigma > 0$, and if $\sigma > 1$ (as is usually estimated) then utility will always be negative!!!!

Don't freak out, if you want to you can always add a constant and make it usually positive. Personally, I think this is completely appropriate. Isn't life nothing more than a constant struggle to get nothing?¹

It is a fine utility function for every $\sigma > 0$, notice that $u'(W) = W^{-\sigma} = \frac{1}{W^\sigma} > 0$ and $u''(W) = -\sigma W^{-\sigma-1} = -\frac{1}{W^{\sigma+1}}\sigma < 0$ so there is absolutely no problem here. And furthermore:

$$r(W) = -\frac{u''(W)}{u'(W)} = -\frac{\left(-\frac{1}{W^{\sigma+1}}\sigma\right)}{\frac{1}{W^\sigma}} = \frac{1}{W}\sigma.$$

This is not only positive but decreasing in wealth, as we suggest it should be.

Now that we have found one utility function with decreasing absolute risk aversion and another with increasing absolute risk aversion then we should also be able to find one with constant absolute risk aversion. The CARA utility function is $u(W) = -e^{-\rho W}$. For this utility function:

$$r(W) = -\frac{-\rho^2 e^{-\rho W}}{\rho e^{-\rho W}} = \rho.$$

Unfortunately this utility function is not too interesting since we don't think that absolute risk aversion should be constant.

¹In case you don't get the joke, let's consider $\sigma = 4$. Then the utility function is:

$$u(W) = -\frac{1}{3W^3}.$$

For any $W > 0$ $u(W) < 0$ but it is increasing in W and as $W \rightarrow \infty$ $u(W) \rightarrow 0$. Get it? You want $W \rightarrow \infty$ but then $u(W) \rightarrow 0 = \text{Nothing??}$ I bet you're rolling on the floor laughing. Who says jokes aren't funny if you have to explain them!!!!

(Sigh, I can't insert a footnote to a footnote in my program. Sigh, such problems. But I want to ask, do you know who says that five exclamation points are the sign of a disturbed mind?)

((And consider this a footnote to a footnote to a footnote. What does it say about someone that they know who says five exclamation points are the sign of a disturbed mind, and then uses them!!!!))

2.5 Relative risk aversion.

To find a measure that quite reasonably should be constant, we define relative risk aversion as:

$$rr(W) = -\frac{Wu''(W)}{u'(W)} = Wr(W) .$$

You might look at this and wonder why we bothered. Well, because constants are nice and the first measure is not reasonably constant. Where did it come from? Lotteries over fractions of wealth ($\sigma \varepsilon I$) and fees that are fractions of your wealth (PI). What's that? You want to seem me derive it? Why sure! Just a second...

(Insert ten pages of math drivel.)

Well, wasn't that exciting? Now, what is this measure for several standard utility functions? Quadratic $rr(W) = \frac{2\beta W}{\alpha - 2\beta W}$, CARA $rr(W) = \rho W$, CES $rr(W) = \sigma$. We can now call CES by its standard name in this environment. Constant relative risk aversion, or CRRA utility. The usual estimates of σ are between three and five.

I should mention that CES gives a constant for risk aversion **only** when we do not work. Models of Labor usually start with the simple alternative of choosing between consumption (c) or leisure (l). You have the constraint that your labor (L) plus your leisure make up all the time you have, or $l + L = T$. We normalize the price of consumption to one, thus the budget constraint is now $c \leq W_0 + wL$, where W_0 is your initial wealth level and w is the wage rate. With CES preferences:

$$u(c, l) = \frac{c^{1-\sigma}}{1-\sigma} + \chi \frac{l^{1-\sigma}}{1-\sigma}$$

the coefficient of absolute risk aversion is:

$$r(W_0) = \frac{\sigma}{c(w, W_0)}$$

where $c(w, W_0)$ is the optimal demand for consumption. This works out to:

$$r(W_0) = \frac{1}{W_0 + Tw} \sigma \left(\frac{\chi^{\frac{1}{\sigma}}}{w^{\frac{1}{\sigma}(1-\sigma)}} + 1 \right) .$$

Which is not a constant, it is in fact affected by the wage rate in two ways. First the wage rate affects the *potential* wealth, $W_p = W_0 + wT$, second it affects the share of potential wealth spent on consumption rather than leisure $\left(\frac{\chi^{\frac{1}{\sigma}}}{w^{\frac{1}{\sigma}(1-\sigma)}} + 1 \right)$. If $\sigma = 1$ (Cobb-Douglas or log preferences) then the second effect is not present, and $r(W_p)$ depends on how much people value leisure (χ). This is disturbing, notice that for any σ if someone likes leisure more (χ high) it makes that person more risk averse ($r(W)$ increases.)

3 Applications

3.1 On why Firms should be Risk Neutral—The Economies of Scale of uncertainty.

A very simple fact is that if someone is involved in n different (partially independent) interactions then they face less uncertainty than does someone who is only involved in one. This translates to firms being less risk averse than consumers, and we simply assume they are risk neutral—or they maximize expected profits. Anytime one party is less risk averse than another there is a great potential for mutually beneficial trade between the two. The less risk averse agent (the firm) can make a profit by absorbing risk.

Say that an insurance firm insures one car for 1000 TL/year. They know that usually administrative and repairs cost for a car is 100 TL, but there is a 5% probability of it being 1000 TL or more. Now, what if they insure two cars? Then they only have to pay more than they took in with .25% probability—because each car has to independently cost them more than 1000 TL. Indeed, you should all know that if the variance of one observation is σ^2 then the variance of the average of n observations is σ^2/n . What you don't think about is that that means that the variance of the total is σ^2 —which as n goes to infinity becomes a smaller and smaller fraction of the total output. The probability density function will also become more and more concentrated, the peaks will get higher and the valleys lower until essentially all uncertainty disappears. All of these things boil down to one simple fact: if a firm sells many units the amount of uncertainty they face goes to zero.

As an example say I know there is a 1% chance of a can opener breaking. As a consumer this would be a big deal, because I only have one can opener. As a producer of thousands of can openers, I translate that into the simple "fact" that 10 out of each thousand will break. I know there will be some variance on that, but so little I can basically ignore it.

The firm that got car insurance from probably insures thousands or millions of cars. Thus they know that while you *might* have a car accident this year they can be sure that $x\%$ of the cars they insure will, and they will have to shell out around y TL per car. You live in an uncertain world, but they essentially live in a certain world.

This is why insurance companies need to be large, the larger they are the less risk they face, and also why they will be risk-neutral. Being risk neutral means that they simply maximize their expected profit. And allows for a great many beneficial trades.

Remark 11 *A risk neutral person can offer a risk averse person a contract where they absorb risk, this contract will benefit both parties.*

3.2 Insurance

The simplest and most obvious application of this is an extremely simple insurance contract. Say that with probability $\tau > 0$ you lose $x > 0$ TL. Alterna-

tively, you can buy an insurance contract for f , when will you buy this?

You will buy this if:

$$(1 - \tau) u(w) + \tau u(w - x) \leq u(w - f) .$$

The left hand side is your expected utility if you don't take the insurance, the right hand side is your expected utility if you do. We know that $f \geq \tau x$ is necessary for the firm to make profits, so is a contract possible? Because you are strictly risk averse the answer is absolutely yes. By strict concavity (risk aversion) we know that:

$$(1 - \tau) u(w) + \tau u(w - x) < u((1 - \tau)(w) + \tau(w - x)) = u(w - x\tau)$$

so if $f = \tau x$ you will be strictly better off taking the policy. This means the firm can charge some fee that is strictly higher than τx , make a little money, pay for their administrative costs and both consumers and the firm will be strictly better off. They can profit from removing risk from your life, both of you like that.

3.3 Return Policies

When I first came to Turkey only international firms would allow you to return your goods, now this has spread to basically all firms—only very small firms don't offer return policies. In Germany some firms allow you to return clothes up to a year after the purchase date (maybe longer, I'm sure of at least a year). In the United States a special exception has to be made for ball gowns. High school and college girls will often buy a ball gown for one night, and then try to return it. So if you buy a ball gown you might have to return it with all the tags still attached—but they don't cancel the general policy, just make an exception. Why? Why is it so profitable to have these return policies?

It's actually fairly simple to analyze. Recognize that with probability $\tau > 0$ the good will not give you any value, assume otherwise that it's value is v . Without a return policy the purchase decision is to buy at the price p if:

$$(1 - \tau) u(w + v - p) + \tau u(w - p) \geq u(w) . \tag{1}$$

Since you are strictly risk averse, we know that:

$$\begin{aligned} (1 - \tau) u(w + v - p) + \tau u(w - p) &< u((1 - \tau)(w + v - p) + \tau(w - p)) \\ &= u(w + (1 - \tau)v - p) \end{aligned}$$

so we must have

$$u(w + (1 - \tau)v - p) > u(w)$$

and since this utility function is strictly increasing ($u' > 0$) this means:

$$w + (1 - \tau)v - p > w$$

and the highest price they can charge is some $p < (1 - \tau)v$. On the other hand if there is a return policy the decision to buy at a price P is:

$$\begin{aligned} (1 - \tau)u(w + v - P) + \tau u(w) &\geq u(w) \\ (1 - \tau)u(w + v - P) &\geq (1 - \tau)u(w) \\ w + v - P &\geq w \\ v &\geq P. \end{aligned}$$

Thus you will buy assuming you will benefit from the good. You can readily see that if $p = v$ in equation 1 you will not buy. In fact, since $p < (1 - \tau)v$, we know that the firm can get a strictly higher revenue by offering a return policy.

The best bit of this policy is that often we are also more willing to buy, knowing that if it doesn't do what we want we can return it. Indeed I myself have often bought thinking this and then never bothered to return it. Return policies are spreading world wide because they are strictly profitable. There is money to be made absorbing the risk consumers face.

3.4 Free Parking

I wrote the article, so you have to learn about it.² Another simple application of this insight is free parking. When you go shopping you know that sometimes, with probability τ , you don't find what you want. If you do, you have to pay a price p and get a value v . However if you pay for parking (a fee of $f > 0$) you pay before you know whether you value the good, so the shopping decision is:

$$(1 - \tau)u(w + (v - p) - f) + \tau u(w - f) \geq u(w) .$$

As usual we know:

$$\begin{aligned} (1 - \tau)u(w + (v - p) - f) + \tau u(w - f) &< u((1 - \tau)(w + v - p - f) + \tau(w - f)) \\ &= u(w + (1 - \tau)v - ((1 - \tau)p + f)) . \end{aligned}$$

This means that $(1 - \tau)v > (1 - \tau)p + f$ and the right hand is their total expected revenue. On the other hand if parking is free they can set $p = v$ and have a strictly higher total revenue.

Thus, like a return policy, free parking is a type of insurance which encourages you to go to their store. You expect prices to be higher as a result, but you are willing to pay the expense. This incentive is so strong that when the Istanbul government declared that parking in shopping malls had to have a fee many malls refused to charge for parking and just paid the fine.

3.5 Wages versus Bonuses

Another strange but simple application is wages. Why do firms give employee's wages instead of simply giving them a fraction of the profit from their activities?

²Hasker, K. and E. Inci. (2014) "Free parking for all in shopping malls." *International Economic Review*, 55:1281-1304.

Some firms (Law firms) do this, but most workers get paid a fixed wage based on the hours they worked. Higher level employees are sometimes given bonuses, but these bonuses are usually a small fraction of their total income. It is a "bonus" after all, not your "income." One of my friends recently changed from a sales job—where nearly one hundred percent of his income was based on his sales—to a salaried job in the same company. He did this despite the salary being below his average income from sales. Why?

The reason in all of these cases is that profit is going to be highly stochastic, uncertain. It might be high or low depending on how the company is doing. For a standard worker a lot of that uncertainty will be outside her or his control, and they are willing to forgo average income for a stable salary.

Thus the firm's owners—who hopefully have many employees and thus face less uncertainty—get the profits, and part of their costs are the wages they have to pay their workers. Among other things, this explains why business owners on average get a higher wage than their employees. It's not because they are more important, it's because they have to deal with all the uncertainty of bad years. They absorb risk, so on average they have to be getting a higher wage.

3.6 A Couple of Common English Sayings and how they apply to risk.

3.6.1 "Don't put all your eggs in one basket" and Diversification

A very common English saying is "Don't put all your eggs in one basket." This insight is soundly based on risk aversion. If you have your eggs in multiple baskets (you went shopping with your kids and each person has their own basket) then the chance of all of them getting broken is close to zero. While you, of course, may trip and drop your basket. Of course the area that we most commonly think of this is investment. You need to diversify your portfolio, put your investments into various different kinds of companies **and** different companies that are the same.

I had a friend who received stocks as part of his salary. Because he believed in the company he never sold any of those stocks. Then the company went bankrupt. Not only did he lose his job but a major chunk of his investments were also now worthless. Trust me when I say that there was no reason to expect the company was going bankrupt, but still he committed a rookie investor's error. Don't put all your eggs in one basket. Even if you love and believe in your company invest in other firms. In fact it would be great to invest in your company's direct competitors.

To see this we are going to assume that there are two stocks you can invest in with independent returns and the same average return, \bar{R} . Their returns also have the same variance, σ^2 , and we assume that the returns are normal.

If the investor put's $\lambda \in [0, 1]$ into the first stock his expected return will be

$$\bar{R}(\lambda) = \lambda\bar{R} + (1 - \lambda)\bar{R} = \bar{R}$$

and so he is indifferent on this level, but what about the variance of his returns? I won't go through the math to show it but:

$$\sigma^2(\lambda) = \lambda^2\sigma^2 + (1-\lambda)^2\sigma^2 = \sigma^2(2\lambda^2 - 2\lambda + 1)$$

So since she is risk averse she dislikes variance, and she should choose λ to minimize the variance. Or

$$\frac{d\sigma^2(\lambda)}{d\lambda} = \sigma^2(4\lambda - 2), \lambda = \frac{1}{2}$$

Of course in the real world it is more complicated than that. You want some "sunny weather" stocks and some "rainy weather" stocks. You might, for example, want to invest in real estate because that's always safe. I actually did this, and then... the 2008 real estate crisis took down a great deal of the world's economy.

Notice that what we just figured out is very similar to our explanation of why firms are risk neutral. In both cases increasing the number decreases the uncertainty.

3.6.2 "Don't put off to tomorrow what you can do today" and Flexibility

I am really torn about this one. After all I am actively encouraging you to study each and every week through my quizzes. Let me just stress this insight is about uncertain situations, and all the uncertainty in my class ended on the first day. Thus you should follow this advice in my class, but in situations of uncertainty... well... don't.

The key thing to recognize is that as time goes on uncertainty will be more and more resolved. With uncertainty decreasing you can make better decisions about what to do. Thus if new information is coming in it is better to delay. Flexibility is always best in an uncertain world. You *should* put off until tomorrow what you can do today if you will have more information.

However this is common knowledge, thus prices will reflect this and often large amounts of money can be made by guessing right. For example stock prices often move before some major announcement about a firm. The traders are trading basing on guesses about what the results will be, and once the announcement is made those who guessed right will make large amounts of money. This also explain why a report—about GDP for example—can be positive but the stock market reacts negatively. The average investor expected the number to be even higher than it was.

A less volatile example is how firms building apartment buildings will generally offer you a deep discount if you buy before the project is finished. This discount is due to a lot of the value of the apartment depending on the actual market at the time the building is finished, since that is unknown they offer you a lower price.

So it is optimal to delay, but often times the market will pay you for the risk you are absorbing by not delaying.

This should make clear the economic value of information. Information is a valuable good because it reduces uncertainty.