Prof. Jacobson's Notes on Risk and Uncertainty

Risk and uncertainty are things I'm really interested in. I study risk preferences in some of my research. Therefore I wrote this handout to organize and summarize the ideas in this topic I think are most important for students.

A lot of the work on how people behave when facing risk and uncertainty comes out of microeconomic theory and laboratory experiments. (There's also some work using "observational data," e.g. people's choices in insurance markets.) Much of this is by economists, but there's increasing cross-disciplinary work, with the closest complement obviously being psychology. In this document I'll focus on theory and practice from a modern microeconomic standpoint.

Risk and Uncertainty

We use the words "risk" and "uncertainty" to refer to situations in which a person doesn't know for sure what the outcome of an event will be. That event could be the flip of a coin, the weather, the outcome of a baseball game, or whether a job offer comes through. Economists refer to uncertain things as bets, gambles, lotteries, and risky prospects, pretty much interchangeably. From the point of view of the person making a decision, we model these events as stochastic—that is, there is random variation that drives what will happen. Another way we describe this is to say there are multiple possible future "states of the world"—such as the state in which it's rainy and the state in which it's not—and we don't know which will prevail. We can think of a few relevant categories.

First, there is "*risk*," also known as "*uncertainty*." In risk, a person knows that there is a certain set of possible outcomes, and the person knows the probability with which each outcome will happen. A classic example of this is a coin flip. Lotteries and slot machine gambles are like this as well. Each risky event with K possible outcomes (states of the world) is fully represented by the list of outcomes $(x_1, x_2, \ldots, x_K)$ and the corresponding probabilities $(p_1, p_2, \ldots, p_K)$.

Next, confusingly, there is "*uncertainty*," also known as "*ambiguity*" or "true" or "Knightian" uncertainty, after Frank Knight. This is the more common case in the real world, because in these cases we know the possible outcomes but we don't know the probabilities with which each one occurs. This is the case for climate change's damages, for example—we can list out a bunch of scenarios but we really can only guess at probabilities. We usually assume that people act as if they make up subjective probabilities—"guesstimates" of the likelihood of each event. In reality, people do seem averse to this situation—they would prefer to know probabilities than to not know them. (If you're interested in this, look for the Wikipedia page on the "Ellsberg paradox.")

Another kind of risk is *strategic risk*. Ordinary risk pits you up against some random process that happens in the real world. In strategic risk, the uncertainty comes from the fact that you don't know for sure what another person is going to do. This could occur in any of a few ways. The person may be playing a mixed strategy, e.g. in a game like rock-paper-scissors. On the other hand, maybe there are different types of people who behave in different ways. Imagine you're trying to decide whether to trust a new friend with a secret. Say there are two kinds of people in the world: trustworthy and untrustworthy people. You might even know what percent of the population is made up of each type. But you don't yet know whether your new friend is a trustworthy type. Whether you share your secret depends on your assessment of strategic risk.

Much of economic modeling focuses on risk (rather than Knightian uncertainty), and we often model both strategic and ordinary risk identically.

Expected Value

Before we define risk preferences, we need a fundamental concept. *Expected value* (EV) is essentially the average outcome you'd expect to happen if this random / risky / stochastic event was played over and over many times. Mathematically, it's the probability-weighted sum of the possible outcomes:

$$EV = p_1 x_1 + p_2 x_2 + \ldots + p_K x_K$$

Risk Preferences

Different people have different tastes (or preferences, or tolerances) for risk. We categorize people based on these tastes. Before I go into those categories, let me make two points.

First, it's fair to assume that people consider "risk" and "reward" when facing uncertain outcomes. But people should consider both how big the possible gains or losses are *and* the probabilities according to which they will get those gains or losses. When faced with a choice among different events that are uncertain, each of us should prefer the one that has higher payoffs if the probabilities are the same… and each of us should prefer the one that has probabilities weighted more toward the best outcomes if the payoffs are the same. On a related note, we will sometimes distinguish between *upside risk* and *downside risk*. Conversationally, people often use the word risk to refer just to downside risk, but in economics risk just means variance. So we want to be exposed to as much upside risk (possibility for good outcomes) as possible while minimizing downside risk (possibility for bad outcomes)

Second, let's define a term. If a risky event is *actuarially fair*, that means one of two things. One is that the average outcome of the event is zero—imagine a coin flip in which you win $1 on heads or lose $1 on tails. The other thing it can mean comes up when there's a cost you can bear to increase or decrease a risk—the price of a lottery ticket, the cost of insurance (more on

insurance shortly), etc. For actuarially fair events, the cost to enter into the "gamble" is the same as the expected value of the gamble. That is, say that there's a coin flip bet where you win $5 on heads and lose $1 on tails. It's an actuarially fair bet if someone charges you the average outcome (which is $2) to play. Similarly, you are paying an actuarially fair amount for insurance if the amount you pay is equal to your expected losses (again, more on this later).

The simplest sort of taste for risk is *risk neutrality*. Risk neutral people care only about expected value. They care don't care about variance. A win $1 / lose $1 coin flip bet is the same to a risk neutral person as a win $1,000 / lose $1,000 coin flip bet. They will choose to enter a bet in which they would win $1,000,001 or lose $1,000,000 with equal probability. A risk neutral person is indifferent about actuarially fair bets and will take actuarially favorable bets.

However, not everyone is risk neutral. In fact, the evidence is strong that most people are at least somewhat *risk averse*: they prefer a higher expected value, but they also don't like facing variance. They're willing to trade off expected value against variance—they will choose an option that gives them a somewhat lower expected value if that reduces the variance they face. This is not irrational—this is just a preference against uncertainty. Risk averse people will never choose to take an actuarially fair bet, and will avoid some actuarially favorable bets if the higher expected value isn't enough to compensate for the variance.

Some people, on the other hand, may be *risk loving*. Risk loving people like a higher expected value, but they also like a higher variance. You can think of these as "thrill seekers," but this again is not irrational—it's just a different preference. Risk loving people will choose to enter actuarially fair bets, and they'll even take some actuarially unfair bets.

Another phrase you may have heard and read is "*loss aversion*." This is a distinctly different concept. It comes out of *prospect theory* and cumulative prospect theory, proposed originally by Kahneman and Tversky. Loss aversion simply means that if you are looking at a gain and a loss of equal size, the loss "looks" bigger to you. The loss makes you sadder than a gain of the same size makes you happy. There is some evidence of this phenomenon in people's choice behavior. We won't go into prospect theory and its extensions in this class, but it's interesting to read about, so I encourage you to look into it if you're interested in these "behavioral" elements.

Risky Choice

So how do people make decisions when the decision involves uncertainty? Assume a person can take one of at least two possible actions. At least one of those actions culminates in an uncertain payoff for the person. For example, in the morning I can choose the actions (Bring Umbrella) and (Leave Umbrella at Home). Bringing the Umbrella causes me to lose some utility regardless of what happens. If I Leave the Umbrella at Home, on the other hand, there's a chance (a probability I will ascertain by looking at the weather forecast) that I will get drenched in a

downpour. Which action do I prefer: the "sure thing" loss associated with Bringing the Umbrella, or the gamble associated with Leaving the Umbrella at Home?

The simplest model of risky choice is that people always choose the option that gives the higher expected value. This model perfectly describes the way that risk neutral people choose.

What about risk averse people? Even if you are averse to risk, is it rational to give up some expected payoff to avoid risk?

You might argue that (at least for small-stakes gambles—that is, when the possible payoffs are small relative to some standard like your annual earnings) it is not rational to be risk-averse because there are ways you can mitigate risk. One way is *diversification*. When we are diversified, we face many risks but we don't expect them to all go bad at the same time. Consider putting together an investment portfolio. Bonds do well in bad times and stocks do well in good times. You want to own both stocks and bonds so that your overall portfolio performance is subject to less variance that it would face if you owned only one or the other. Think also of lobbyists: their safest bet is to back both the Republican and Democratic candidates in a race. The value of each investment is uncertain, because you don't know which one will win. But if you back both of them, you know that one of your bets will pay off.

Diversification can reduce risk if the uncertainties in the different risky assets are uncorrelated with each other or if they are negatively correlated. An example of uncorrelated risks may be two different stocks from different sectors on the stock market—say, General Motors and Apple. (They're not perfectly uncorrelated, but pretend for the moment that they are.) Mutual funds consist of lots of individual stocks so investors aren't fully exposed to the risk associated with any one stock—at any given time, some stocks will go down but others will go up, so your potential losses are reduced. An example of negatively correlated risks is the Republican and Democratic candidate example from the previous paragraph.

Diversification does not work if risks are positively correlated with each other. A great example is the risks associated with house mortgage default in the run-up to the 2007-8 financial market crash. Investment firms bundled together risky mortgages to create assets they thought were diversified and therefore safe, because while John and Mary each have a chance of default, they probably won't both default at the same time… or so went the logic. As it turned out, lots of John's and Mary's simultaneously defaulted because there were correlated factors driving the defaults, especially once the default wave picked up speed—the bad economy caused lots of job losses, for example, which increased defaults across the board.

At any rate, diversification is an important strategy. It's argued that because of diversification, firms should make individual decisions in a risk averse way. Each individual decision a firm makes likely only gives payoffs that are a small percent of that firm's total revenue stream (or corporate value, etc.). The outcomes of these varied risky choices are likely to be only weakly correlated. Therefore, firms (it is argued) should see each choice as a piece of a portfolio—and

shouldn't sacrifice a higher expected value to try to reduce variance in each little piece. A similar argument can be made for individual people! We make millions of decisions over our lives; when facing a small-stakes risk, maybe we should simply go for higher expected value.

Expected Utility

However, risk cannot be fully mitigated, and in reality people are not risk neutral. The dominant paradigm in modeling risky choice is *expected utility (EU) theory*. This theory allows for risk aversion and risk lovingness.

In expected utility, we assume that a person's value for a risky asset is the probability-weighted sum of *utilities* of outcomes. This may sound like a small difference from EV, but it allows the curvature of the utility function to come into play. This curvature makes it so that people are more or less swayed by the reward from a large prize, and therefore that they are more risk loving or more risk averse respectively.

We now define the utility function of a person as having only one argument: money (or income, or wealth). Obviously, this money represents the ability to buy the utility-conferring goods. Mathematically, a person's expected utility of experiencing a risky event is:

$$EU = p_1 u(x_1) + p_2 u(x_2) + \ldots + p_K u(x_K) \quad [1]$$

A risk neutral person can be represented with a simple utility function that directly maps money into utility in a linear way, such as $u(x) = x$. I will show this in a moment.

Consider a bet in which you could win \$100 if a coin comes up heads, or \$0 if it comes up tails. Let's see the three kinds of people.

First, a risk neutral person will look at only the expected value of the bet:

$$EV = \frac{1}{2}\$100 + \frac{1}{2}\$0 = \$50$$

If we use the simple utility function $u(x) = x$, then expected utility is the same:

$$EU = \frac{1}{2}\$100 + \frac{1}{2}\$0 = \$50$$

Let's plot out the utility this person receives against money. We can represent the bet on these axes as well. Expected utility is the probability-weighted sum of the utilities of the outcomes. We

---

[1] You may be saying to yourself: I thought utility was an ordinal concept, but now it's getting to look pretty cardinal. If so, you're right! Once we get into the realm of expected utility, we are starting to put more restrictions on the utility function and it's no longer purely ordinal.

can therefore draw a line between these two outcomes and the expected utility will be the vertical "height" of a point somewhere along this line. Exactly where along this line depends on those probabilities—if it's a 50-50 bet, the point will be the halfway point, whereas if the probability favors one outcome the point will be closer to that outcome. All of this is based on the logic that the "average" (or a "moderate combination," if you remember back to that from consumer choice theory) of two points is on the line between them.
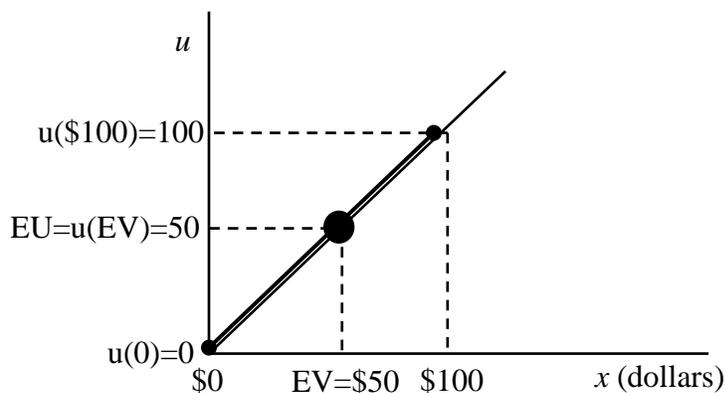

Now, let's consider a risk averse person. A risk averse person will have a concave utility function, such as $u(x) = \sqrt{x}$. For this person, the utility of getting the expected value of the bet for sure is much higher than is the utility of entering the bet:

$$u(50) = \sqrt{50} = 5\sqrt{2} \cong 7.07$$

$$EU = \frac{1}{2}\sqrt{100} + \frac{1}{2}\sqrt{0} = 5$$

Now that the utility function is concave, the chord connecting the two lottery outcomes falls below the utility function. Thus the weighted sum of the utilities of the outcomes is less than the utility of the "average" or expected outcome. A risk averse person always prefers to get the expected value of a bet for sure than to actually have to face the bet with its variance.
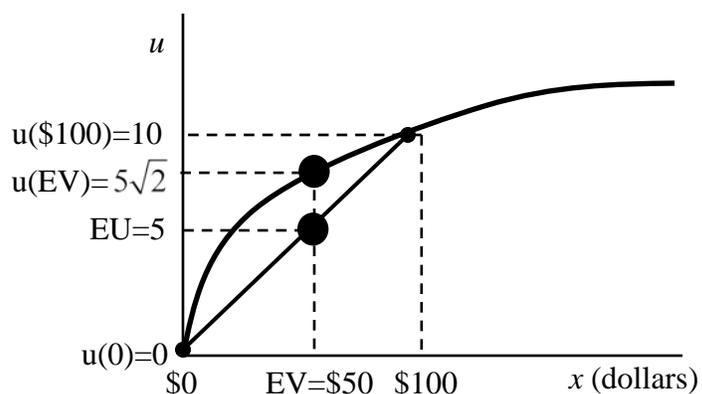
Finally, consider a risk-loving person. A risk-loving person's utility function is convex; for example, a utility function like $u(x) = x^2$. For such a person, the utility of the expected value of the bet is less than the utility of taking the bet, because this person likes to face variance.

$$u(50) = 50^2 = 2,500$$

$$EU = \frac{1}{2}100^2 + \frac{1}{2}0^2 = 5,000$$

Graphically, the chord between the two possible outcomes lies above the utility function, so the expected utility of the bet is above the utility of the expected value. A risk loving person always prefers to play a gamble rather than receive that gamble's expected value for sure.
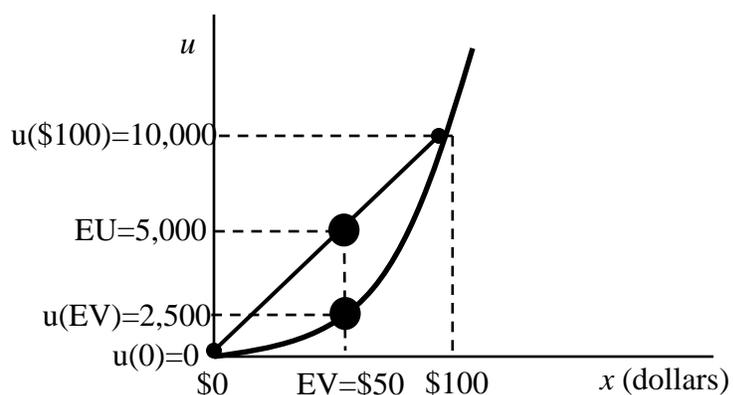
The expected utility model does have some shortcomings. It does not allow for loss aversion, which may be a real phenomenon. And it seems behaviorally that people give too much attention

to low probabilities and too little attention to large probabilities. Both of these are addressed in versions of *prospect theory*. And there are other conceptual and behavioral issues as well. But nonetheless, EU is a reasonable place to start.

Insurance

When someone buys insurance, they are paying money to reduce the risk they face by paying the insurance company to bear that risk for them. You pay a small payment called a *premium* on a regular basis (or in some cases, such as for electronics purchases, you may simply pay an up front amount). The insurance company then promises that if something bad happens, they will bear the cost for you. For example, I pay a monthly premium for health insurance. For this payment, the insurance company provides me the service of paying for large unexpected health bills. The chances I will get hit by a bus are low but it would cost me a lot if it happened; now the insurance company will pay those costs for me if it happens.

Insurance does not save the average person money. Why? The insurance company can't lose money on the average customer, or it will go out of business! The insurance company looks at each of us as a gamble: most likely we'll cost a small amount, but there's a chance we'll cost a lot of money. How much money does the company need to be paid (in premiums) to be willing to be liable for such a large cost?

Insurance companies hire actuaries to figure out the likely risks facing whatever they are insuring. They want to know what a woman of my age who lives in Massachusetts is like—with what probability I will incur which different kinds of health-related costs. Then they want to charge me a rate that is at least as large as the money they expect to lose on me. Actuarially fair insurance is insurance for which the premiums equal the expected payouts. In reality, insurance companies have to charge more than the actuarially fair price of insurance because there are other costs involved in providing insurance. For example, they need to pay the salaries of the agents, and they need to hire the actuaries, and they need to rent a building, etc.

So imagine that there's only one terrible thing that can happen to compromise my health, and that it happens with probability $p = 0.1\%$ each year. If it happens, the costs are $x = \$100,000$. The insurance company will charge me an annual premium greater than $px = \$100$—say, $150—to insure me. If nothing happens, I will have paid that amount and the insurance company keeps the money. If the bad thing happens, I will have paid that amount but the insurance company has to pony up $100,000 to cover the medical costs.

Remember, the insurance company insures lots of people. If the risks that hit individuals are not correlated with each other, then the insurance company is diversified. The company should be risk neutral because the chances of all of us needing double bypass surgery are very low.

What kind of people buy insurance, and is it a good idea to buy it? Assume for the moment that insurance companies know the kinds of risks people face—I'll address the issues of adverse selection and moral hazard in a moment. Those issues aside, the only people who would buy insurance are people who are risk averse. When you buy insurance, you're giving up some expected value to reduce your variance. Is it still rational to buy insurance, then? Yes, as long as the premium isn't too high and the potential payout too low. Many of the losses we insure against are large and could wipe out a regular family budget.

Some kinds of insurance, however, are generally bad deals. Insurance to cover small regular expenses is generally a bad deal. An example of this case is eye care insurance (which generally covers a regular eye checkup and new lenses / glasses—these are predictable expenses, and the insurance premium generally costs at least as much as the cost you'd incur). (Some kinds of insurance can use bargaining power to drive down the costs of provision, though, and that can save the consumer money, but that's not a feature of insurance per se.) On the other hand, insurance for most consumer products, such as electronics, is priced at an exorbitant rate—given what can be covered, the probability of having a covered event would have to be ridiculously high to warrant buying the insurance (or you'd have to be insanely risk averse).

Now, adverse selection and moral hazard afflict insurance markets. They are both issues related to information that one party in the transaction has, and both can cause market failures even in a perfectly competitive market without classic externalities. To be more concrete, I'll discuss both in the context of health insurance, although they have applications in many settings.

Adverse selection occurs when there are different types of people and the insurance company can't tell which type any individual person is. Say some people have a high risk of heart disease. Each person knows his own heart disease risk because he knows his family history and his lifestyle. But the insurance company doesn't know which are the risky people. So the insurance company has to set insurance premiums for everyone such that it can cover its costs on average. The premium will be somewhere between the expected cost of insuring a high-risk person and the expected cost of insuring a low-risk person.

If people were risk neutral or not very risk averse, then you'd have a problem. The low-risk people would not want to buy this insurance! It would be a bad deal for them on average, because they'd be paying out more in premiums than they expect to get back in coverage. The low-risk people would then exit the market—they refuse to buy insurance. All that are left in the insurance market are high-risk people. The insurance company then has to raise its premiums to reflect the fact that it knows that low-risk people won't stay in such a market.

This problem is somewhat mitigated if people are relatively risk averse, but even so, if costs of caring for the two kinds of people are high enough then you will have the healthiest people refusing to buy insurance, and the problem cascading from there. This is called a "*death spiral*"

because if there is a continuum of different levels of risk, then you'll have everyone but the very riskiest person leaving the market so that the market basically collapses.

This is why economists have long argued that insurance programs must have an individual mandate—a requirement that everyone buy insurance. This makes it so that the pool of insured people stays strong and healthy, rather than disintegrating so that only the sickest people are buying insurance. As a side note, this idea was originally embraced by politicians from the right, because of the argument about how markets work. How things change! This is now a feature of the Affordable Care Act (Obamacare) that they love to hate.

The other issue is moral hazard. *Moral hazard* occurs when one party (called the principal) can't monitor what the other party (the agent) does, and the agent's actions can be costly to the principal. If the principal is bearing risk for the agent, the agent may take riskier actions knowing that that risk is not totally on his shoulders. In this context, the problem is that the insurance company can't monitor the health behaviors of the insuree. If I did not have health insurance, I might be very careful about risks to my health—skydiving, motorcycle riding, and high impact sports would probably be out, because I would know that if something went wrong I would have to pay the price myself. Since I do have health insurance, I might not be so careful, knowing that at least the monetary costs are not all on me. I might choose to smoke, drink, and engage in all sorts of risky behaviors. In the model, this is a case in which the insuree's actions affect p.

Moral hazard was first discussed in the context of life insurance for babies. People worried that if we allowed life insurance companies to insure the lives of babies, parents might not take as good care of their babies, because if the baby was lost, at least they'd get a chunk of money coming to them. It's a horrible thought and we'd all like to believe that parents don't think this way, but we do have to watch out for bad incentives.

You can also think about it in terms of the Gaudino option. That acts like grade insurance so that your grade can't drop too low in a Gaudino'd class. But will you work as hard in a class knowing that you have Gaudino insurance as you would if you did not?

With regard to both moral hazard and adverse selection, insurance companies do everything they legally can to screen customers for risks and risky behaviors and to penalize those characteristics and behaviors that are most costly. But many things they may wish to do are illegal in some countries. For example, women use more health care than men, so are more costly to insurers; but in the UK, a law was recently passed to disallow insurers from charging women more solely because of their gender. In the US, the Affordable Care Act tried to limit the practice of charging smokers more for health insurance.