Lecture 14: More Uncertainty

1 Where Are We/Logistics

- Today's my last lecture
- Final homework is up, due next Tuesday night at midnight
- Midterm is next Thursday evening, in Ingraham 19
- Extra office hours next week (TBA)
- Online course evaluations for my half of this class have started –
  you should have received an email about it –
  and will be open until Sunday

- Last time, we began decision-making under uncertainty
- We thought about our choice set as lotteries over different outcomes –
  probability distributions over possible prizes we'd get to consume
- We saw that if preferences over lotteries are complete and transitive,
  and satisfy two more assumptions – continuity and independence –
  then preferences can be represented by an *expected utility* function $U(L) = \sum_i p_i u(x_i)$,
  where a utility level is assigned to each outcome,
  and the utility of a lottery is the expected utility value of the prize

- Thinking of the prizes as different amounts of money,
  u is called the Bernoulli utility function,
  and the whole thing the von Neumann-Morgenstern utility function
- We saw that risk aversion – preferring the expected value of a lottery to the lottery itself –
  is the same as concavity of the Bernoulli utility function,
  and that the Arrow-Pratt coefficient of absolute risk aversion, $A(x) = -\frac{u''(x)}{u'(x)}$,
  gives a measure of "how risk-averse" someone is
- (A decisionmaker with a smaller Arrow-Pratt coefficient at every x
  has higher certain equivalents and lower risk premia for all lotteries,
  and is more likely to prefer a risky lottery to a sure thing;
  and if your Arrow-Pratt coefficient is decreasing in x,
  you effectively become less risk-averse as you get wealthier)
- Today, a related point on comparing risk aversion levels;
  then we'll start comparing lotteries,
  asking when can everyone agree that one lottery is "better" than another lottery;
  and we'll see a few behavioral anomalies regarding choice under uncertainty

- Before we start... any questions?

2 Relative Risk Aversion

- We saw last time that the Coefficient of Absolute Risk Aversion, $-\frac{u''}{u'}$,
  gave a good local measure of how risk-averse you are;
  we motivated it initially by showing that the risk premium for a small gamble,
  or the probability premium for a small gamble,
  are both proportional to it
- We saw that if two decisionmakers are ranked via the Coefficient of Absolute Risk Aversion,
  this tells you a lot about their relative attitudes about risk;
  and we could apply the same idea to one person at different wealth levels
- For the latter –
  when we're comparing how I evaluate risks when I'm richer versus when I'm poorer –
  but we could alternatively think of risks in terms of their size relative to my wealth
- that is, rather than thinking of a lottery as a distribution of absolute outcomes,
  like losing \$50 or winning \$100,
  which you might evaluate differently at different levels of starting wealth,
  we can think of a lottery as a distribution of returns applied to your current wealth,
  like losing a half a percent of your wealth or winning a percent,
  and think about how your preferences over proportional risks like that change with w
- In that formulation, a lottery is a probability distribution of relative changes to your wealth level, or a distribution of returns t with

$$U(F) \quad = \quad \int u(tw)dF(t)$$

- In that case, the relevant measure of risk aversion is the **Arrow-Pratt Coefficient of Relative Risk Aversion**,

$$R(u,x) \quad = \quad -\frac{xu''(x)}{u'(x)}$$

- (We could again calculate your risk premium for a small relative risk,
  or your probability premium for a small relative risk,
  and they would be proportional to R)

- If this function is increasing in x, you have **increasing relative risk aversion**, and you become more averse to proportional risks as your wealth level increases; so if you prefer at 2% gain to a fifty-fifty shot at a 5% gain at one wealth level, you still prefer the sure thing at higher wealth levels; if this is decreasing, you have decreasing relative risk aversion, and you become less averse to proportional risks
- Finally, there's a family of utility functions with constant relative risk aversion:

  $$u(x) = \frac{1}{1-\rho} x^{1-\rho}$$

  for $\rho \geq 0$ and $\rho \neq 1$, and its limit $u(x) = \ln(x)$ as $\rho \to 1$
- (For $\rho \in (0,1)$, this is x^a; with $\rho > 1$, this is $-x^{-a}$, it's sometimes written $\frac{1}{1-\rho}(x^{1-\rho} - 1)$ so its range is positive rather than negative when $\rho > 1$)
- This whole family is called Constant Relative Risk Aversion utility; you can easily calculate for yourself that $R(x) = \rho$, and with these utility functions, you'll maintain the same risk attitude about proportional risks as your wealth changes
- (You'll explore this a bit on the last homework assignment.)

3 Subjective Expected Utility

- I emphasized last time that we were dealing with *objective* probabilities –

  events where we knew the exact probability with which they'll occur

- And I mentioned there's a complementary literature on *subjective* probabilities –

  events where we don't know the exact probability,

  like gambles on a football game or an election or the weather

- Now that we've seen the von Neumann-Morgenstern result about when preferences over objective lotteries can be represented by a utility function of the expected-utility form,

  I want to make the link more explicit

- Leonard Savage, in his 1954 book,[1]

  and Anscombe and Aumann, in a 1963 paper[2]

  established similar results for subjective probability:

  if you have uncertainty without objectively-known probabilities,

  or a mix of some probabilities that are objectively known and some that are not,

  and if your preferences over lotteries satisfy certain axioms,

  then your preferences are again represented by a utility function with the expected-utility form,

  where there are now probabilities (not from any exogenous source, but derived from your preferences over lotteries) that we can interpret as your subjective beliefs about the probability of each event

- The details are a lot messier, but the core result is the same –

  under the right assumptions about your preferences over lotteries,

  we get a "subjective expected utility" representation,

  where the probabilities you assign to each event are implied by your preferences over lotteries

[1] see Savage (1954), *The Foundations of Statistics* ch. 2-3

[2] Anscombe and Aumann (1963), "A Definition of Subjective Probability," *Annals of Mathematical Statistics* 34.1.

4 When is a lottery better than another lottery?

- So far, we've established what it means for a decisionmaker to prefer one lottery to another – when $\int u(x)dF(x) \geq \int u(x)dG(x)$
- And we've established what it means for one decisionmaker to be more risk-averse than another – when $-\frac{u''}{u'} \geq -\frac{v''}{v'}$, or when $u = g \circ v$ and g is concave
- Next, we'll think about what it means for one lottery to be better than another, in the sense that *every* decisionmaker would prefer it

4.1 First-Order Stochastic Dominance

- To start: when would *every* decisionmaker – risk-averse, risk-neutral, or risk-loving, regardless of u – agree that lottery F is better than lottery G?

- We could also phrase this as, how could we unambiguously improve a lottery?
- One thing we could do is to start off with a lottery G,

  but then after we run G and find out what prize you'll get,

  maybe we increase it
- That new composite lottery – where for each prize x that you would have received under G,

  you get a prize under F that's *at least* x –

  seems to pretty clearly be better

- What does this look like in terms of density functions?

  start off with one density function,

  and then change it by shifting some probability mass to the right
- What does this look like in terms of CDF?

  The new CDF is lower starting where you took that probability mass from,

  and "catches up" wherever you added the mass

change in density function

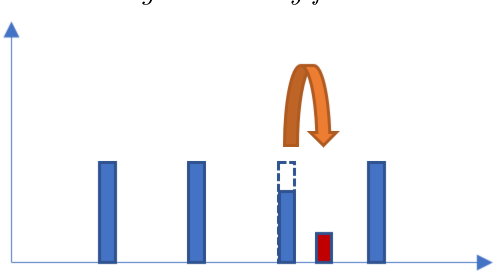

change in CDF

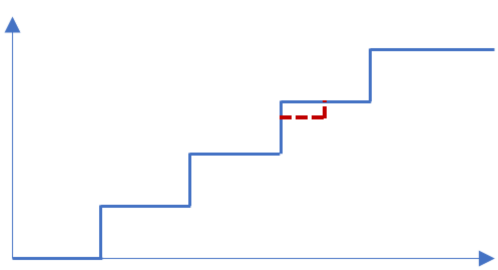

- so, shifting some probability toward higher prizes seems like an uncontroversial way to improve a lottery,

  and seems to relate to lowering the value of the CDF in some places

- It turns out, this gives us a clear characterization of when all decisionmakers should unanimously agree one lottery is better than another

- **Definition.** A lottery F **first-order stochastically dominates** a lottery G, or $F \geq_{FOSD} G$, if $F(x) \leq G(x)$ for every x.

- **Prop.** $F \geq_{FOSD} G$ if and only if $\int u(x)dF(x) \geq \int u(x)dG(x)$ for every nondecreasing u

  (or, $F \geq_{FOSD} G$ if and only if $F \succsim G$ for every expected-utility-maximizer with nondecreasing Bernoulli utility function)

- First, let's prove the "only-if" – that $F \geq_{FOSD} G$ is necessary for this to hold,

  or that if $F \not\geq_{FOSD} G$, there's some expected-utility maximizer who prefers G

- If there's some y such that $F(y) > G(y)$,

  then consider the decisionmaker with Bernoulli utility function

$$u(x) = \begin{cases} 0 & if \quad x \leq y \\ 1 & if \quad x > y \end{cases}$$

- This guy's expected utility from F is $\Pr(x > y) = 1 - F(y)$,

  and his expected utility from G is $1 - G(y)$,

  so he prefers G

- So clearly, if every expected-utility maximizer (for u increasing but not necessarily concave) prefers F to G,

  F must first-order stochastically dominate G, or $F(x) \leq G(x)$ everywhere

- Now let's prove that if $F \geq_{FOSD} G$, everyone agrees it's better
- We'll simplify the proof by assuming F and G have support in $\mathbb{R}_+$, u is differentiable, and $u(0) = 0$, so we can write $u(x) = \int_0^x u'(s)ds$
- Then

$$\begin{aligned}
U(F) - U(G) &= \int_0^\infty u(x)f(x)dx - \int_0^\infty u(x)g(x)dx \\
&= \int_0^\infty \int_0^x u'(s)ds\, f(x)dx - \int_0^\infty \int_0^x u'(s)ds\, g(x)dx \\
&= \int_0^\infty \int_s^\infty f(x)dx\, u'(s)ds - \int_0^\infty \int_s^\infty g(x)dx\, u'(s)ds \\
&= \int_0^\infty (1 - F(s))u'(s)ds - \int_0^\infty (1 - G(s))u'(s)ds \\
&= \int_0^\infty (G(s) - F(s))u'(s)ds
\end{aligned}$$

  so if $F \leq G$ everywhere, since u is increasing (so $u' \geq 0$), the integrand is always positive and $U(F) \geq U(G)$.

- So, if $F(x) \leq G(x)$ everywhere, then every expected-utility maximizer with a nondecreasing Bernoulli function agrees $F \succsim G$

- I'll also note that if $F \geq_{FOSD} G$, then we really can generate F as a composite lottery where first, we run G, and then we give you an outcome of *at least* the outcome G would have given you
- This is because if $F(x) \leq G(x)$ for every x, then $F^{-1}(t) \geq G^{-1}(t)$ for every t; so we can first run G, find out how much money you would have gotten, see what percentile that is of G, and give you the corresponding percentile of the distribution F; and the matching percentile of F is always weakly higher
- So if $F \succsim_{FOSD} G$, we can generate F from G by shifting some probability to the right

4.2 Second-Order Stochastic Dominance

- Next, we'll consider the case where two lotteries have the same expected value,

  but one is unambiguously **more risky** than the other

- For parametric distributions, this may be obvious –

  if we're comparing two normal distributions with the same mean and one has higher variance,

  it's obvious that's the more risky one,

  and anyone who's risk-averse will prefer the lower-variance one

- but we'd like a more general condition that captures the same intuition

- So first, how would we make a distribution more risky?

- This time, let's think about starting with the distribution F,

  learning the outcome x you got under F,

  and then replacing that outcome x with a new lottery that has expected value x

change in density function *change in CDF*

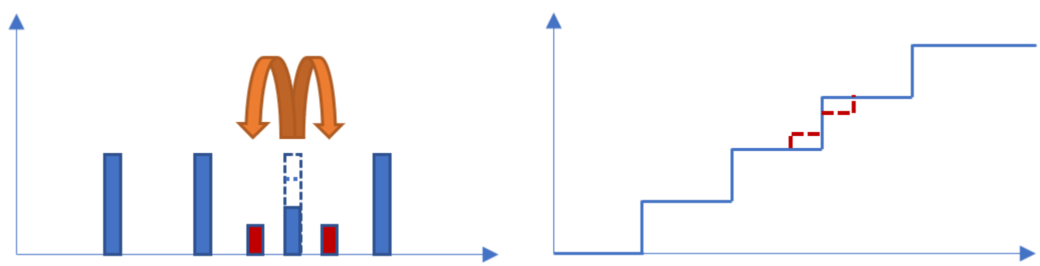

- This is called a "mean preserving spread" –

  you spread out the outcomes of F, without changing the expected value,

  to get the new distribution G

- What does this do to the CDF?

- It increases it to the left, and decreases it to the right,

  so the comparison will not be quite as simple as before

- But we still get a very clean characterization of when one distribution is riskier than another

- **Definition.** F **second-order stochastically dominates** G, or $F \geq_{SOSD} G$, if they have the same expected value and

$$\int_{-\infty}^{x} F(s)ds \quad \leq \quad \int_{-\infty}^{x} G(s)ds$$

  for all x.

  (Note $E_F = E_G$ implies that $\int_{-\infty}^{\infty}(F(s) - G(s))ds = 0$ when taken over the whole support.)

- Just like Pratt showed us the equivalence of several notions of a "more risk averse" decision-maker,

  Rothschild and Stiglitz showed the equivalence of several notions of a riskier *lottery*

- **Theorem (Rothschild and Stiglitz).** Let F and G be two lotteries with the same expected value. The following definitions of "G is more risky than F" are equivalent:

  1. $\int_{-\infty}^{x} F(s)ds \leq \int_{-\infty}^{x} G(s)ds$ for every x
  2. $F \succsim G$ for any risk-averse expected utility maximizer,
     or $\int u(x)dF(x) \geq \int u(x)dG(x)$ for any increasing, concave u
  3. G is derived from F via a mean-preserving spread –
     the distribution of outcomes in G can be derived by first taking a draw from F and then adding mean-0 noise

- (To be formal, if $X \sim F$ and $Y \sim G$, condition three is that $Y \stackrel{d}{=} X + Z$,

  where Z is another random variable which is not necessarily independent of X

  but has $E(Z|X = x) = 0$ for every x, and $\stackrel{d}{=}$ means "has the same distribution as")

- I'm not going to give the full proof
- to see that 1 and 2 are equivalent, in the case where u is twice differentiable, when $E_F = E_G$, we can show[3]

$$U(F) - U(G) = \int_{-\infty}^{\infty} (-u''(x)) \left[\int_{-\infty}^{x} G(s)ds - \int_{-\infty}^{x} F(s)ds \right] dx$$

- So if u is concave, then $-u'' \geq 0$, and so the integral condition implies $U(F) - U(G) \geq 0$

- On the other hand, if the integral condition is violated, then there's some y at which $\int_{-\infty}^{y} G(s)ds < \int_{-\infty}^{y} F(s)ds$ (and since CDFs are always within $[0,1]$, their integrals are continuous, so the same holds in a neighborhood around y); so if we pick u such that $u'' \approx 0$ away from y, and negative very close to y, and get that $U(F) - U(G) < 0$ for that u
- (show what the u looks like – it's basically $u(x) = \min\{x, y\}$, just smoothed a little to be differentiable right around y)

[3]We already showed that $U(F) - U(G) = \int u'(x)(G(x) - F(x))dx$. If we assume u is twice differentiable, and abuse notation a little and write $u'(x) = u'(\infty) - \int_x^{\infty} u''(s)ds$, this is

$$\begin{aligned} U(F) - U(G) &= \int_{-\infty}^{\infty} \left(u'(\infty) - \int_x^{\infty} u''(s)ds \right) (G(x) - F(x))dx \\ &= u'(\infty) \int_{-\infty}^{\infty} (G(x) - F(x))dx + \int \int_x^{\infty} (-u''(s))ds(G(x) - F(x))dx \end{aligned}$$

Integration by parts allows us to show $E_F = E_G \leftrightarrow \int (G(x) - F(x))dx = 0$, so the first term vanishes; switching the order of integration on the second gives

$$U(F) - U(G) = \int_{-\infty}^{\infty} \left(\int_{-\infty}^{s} (G(x) - F(x))dx \right) (-u''(s))ds = \int_{-\infty}^{\infty} \left(\int_{-\infty}^{s} G(x)dx - \int_{-\infty}^{s} F(x)dx \right) (-u''(s))ds$$

- Rothschild and Stiglitz show that the integral condition is also equivalent to being able to get to G from F via a series of mean-preserving spreads,

  by constructing the mean-preserving spreads for discrete distributions and making a limit argument for the continuous case

- so, the integral condition $\int_{-\infty}^{x} F(s)ds \leq \int_{-\infty}^{x} G(s)ds$ holding at every x,

  is equivalent to every risk-averse decisionmaker preferring F to G,

  which is equivalent to "$G = F+$ noise" – you can generate G from F via a series of mean-preserving spreads

4.3 a technical note I'll skip in lecture

I'll skip this in lecture, and feel free to skip over it, but I'm including it here because it's neat. Here's a paragraph from Jonathan Levin's 2006 lecture notes on Choice Under Uncertainty, available online:

"[The FOSD and SOSD proofs] rely heavily on integration by parts. There is, however, a general way to prove stochastic dominance theorems that doesn't require differentiability. The idea is that given a class of utility functions $\mathcal{U}$ (e.g., all non-decreasing functions, all concave functions, etc.), it is often possible to find a smaller set of "basis" functions $\mathcal{B} \subset \mathcal{U}$ such that every function $u \in \mathcal{U}$ can be written as a convex combination of functions in $\mathcal{B}$. (If the set $\mathcal{B}$ is minimal, the elements of $\mathcal{B}$ are "extreme" points of $\mathcal{U}$ – think convex sets.) It is then the case that $\int udG \geq \int udF$ for all $u \in \mathcal{U}$ if and only if $\int udG \geq \int udF$ for all $u \in \mathcal{B}$, so one can focus on the basis functions to establish statistical conditions on G and F. You can try this yourself for first order stochastic dominance: let $\mathcal{B}$ be the set of "step" functions, i.e., functions equal to 0 below some point x and equal to 1 above it. For more, see Gollier (2001) [Christian Gollier (2001), *The Economics of Risk and Time*, MIT Press.]; this was also the topic of Susan Athey's Stanford dissertation."

The idea that "an increasing function is a convex combination of step functions" is clearest when you consider the differentiable case, since (up to an additive constant) we can write differentiable u on $\mathbb{R}_+$ as

$$u(x) \quad = \quad \int_0^x u'(y)dy \quad = \quad \int_0^\infty u'(y)b_y(x)dy$$

where

$$b_y(x) \quad = \quad \begin{cases} 0 & if \quad x < y \\ 1 & if \quad x \geq y \end{cases}$$

So the step functions $\{b_y(x)\}$ serve as a basis for the increasing and differentiable functions on $\mathbb{R}_+$, with $u'(y) \geq 0$ giving the weight assigned to each basis function. Since $\int_{-\infty}^{+\infty} b_y(x)dF(x) = 1 - F(y)$, it's easy to see that $F(\cdot) \leq G(\cdot)$ everywhere if and only if $\int b_y(x)dF(x) \geq \int b_y(x)dG(x)$ for every basis function b_y. And if we write any increasing Bernoulli utility function $u(\cdot)$ as a convex combination of basis functions $b_y(\cdot)$ with weights $w(y)$, then

$$\int u(x)dF(x) \quad = \quad \int \left(\int w(y)b_y(x)dy \right) dF(x) \quad = \int w(y) \left(\int b_y(x)dF(x) \right) dy$$

so if we know $\int b_y(x)dF(x) \geq \int b_y(x)dG(x)$ for each basis function b_y, then it holds for every increasing Bernoulli utility function.

If you want to try this for *second*-order stochastic dominance, then up to an additive constant, a basis for the increasing, *concave* functions is the kinked functions

$$b_y(x) \quad = \quad \min\{x, y\}$$

which are increasing with slope one on $x < y$ and then constant on $x > y$.

5 Behavioral Critiques of Expected Utility Theory

- I mentioned last week that expected utility theory has been widely critiqued by a variety of behavioral economics work
- Here are some classic examples of ways it doesn't seem to match observations

5.1 Risk aversion over small gambles requires "a lot" of concavity

- We saw earlier that for small gambles,

  the risk premium scales with the square of the size of the gamble
- This makes sense, since risk aversion is related to the curvature of the utility function, not the slope
- However, this means that expected utility maximizers should naturally become very nearly risk-neutral when the stakes get small
- Matthew Rabin has a paper[4] showing the flip side of this –

  if someone is an expected-utility maximizer and is fairly risk-averse over modest stakes,

  they must be *extremely* (unrealistically) risk-averse over large stakes
- He uses this to basically argue that expected utility isn't a useful theory at all levels –

  that you can't take seriously its implications on both large risks and small risks,

  if you assume someone has a single, consistent expected utility function
- One example he gives: suppose that at any starting wealth level w,

  you would reject a 50-50 chance to either win \$110 or lose \$100
- If that's the case, and you're an expected utility maximizer,

  Rabin shows that your utility function must be so concave that you would reject a 50-50 shot to either lose \$1000 or win \$10 *billion*
- Rabin therefore argues that if people are risk-averse over small gambles,

  it can't be because their utility function is curved, but because it must be kinked at zero –

  people get more disutility from a loss than utility from a gain, or are "loss-averse"
- (Ariel Rubinstein and others have responded that being that risk-averse "at any starting wealth level" is unrealistic and that's the "problem" here;

  but I think it's valid that concavity of u at the levels we want to assume for reasonable "big-picture" behavior can't generate substantial aversion to very small risky gambles;

  if you don't want to flip a coin to lose \$1 or win \$1.10, it can't be because your u is concave)

[4] M. Rabin (2000), "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica* 68(5).

5.2 Prospect Theory

Daniel Kahneman and Amos Tversky, "Propsect Theory" (1979) – the paper that basically launched behavioral economics – showed a bunch of experiments demonstrating that peoples' choices are often inconsistent with expected utility maximization[5]

People overvalue the "sure thing"

- One example, the Allais Paradox – people seem to "overvalue" certainty
- Given a hypothetical choice between the following two lotteries:

A			B		
\$2500	*w.p.*	.33			
\$2400	*w.p.*	.66	\$2400	*w.p.*	1
\$0	*w.p.*	.01			

  82% of subjects said they'd choose B, the sure thing

  But, given the choice between:

C			D		
\$2500	*w.p.*	.33	\$2400	*w.p.*	.34
\$0	*w.p.*	.67	\$0	*w.p.*	.66

  83% of subjects chose the first

- So at least 65% said $B \succ A$ and $C \succ D$

[5] D. Kahneman and A. Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47(2).

- The Allais Paradox, to me, is really a rejection of the independence assumption
- Start off with the two lotteries

X			Y		
\$2500	$w.p.$	$\frac{33}{34}$	\$2400	$w.p.$	1
\$0	$w.p.$	$\frac{1}{34}$			

- By completeness, either $X \succsim Y$ or $Y \succsim X$.
- The lotteries A and B are mixtures of X or Y, respectively,

  with the sure thing \$2400 (or the lottery Y);
- the lotteries C and D are mixtures of X or Y, respectively,

  with the sure thing 0
- Independence says if $X \succsim Y$, then $A \succsim B$ and $C \succsim D$;

  if $Y \succsim X$, then $B \succsim A$ and $D \succsim C$
- But there are no rational preferences with independence that allow $B \succ A$ and $C \succ D$
- Although that's what people seem to prefer –

  and it seems to be about putting extra value on eliminating risk altogether

- (To see more explicitly that expected utility maximization can't allow $B \succ A$ and $C \succ D$, note that

$$B \succ A \quad \longrightarrow \quad u(w+2400) \quad > \quad .33u(w+2500) + .66u(w+2400) + .01u(w)$$

$$C \succ D \quad \longrightarrow \quad .34u(w+2400) + .66u(w) \quad < \quad .33u(w+2500) + .67u(w)$$

and subtracting the second from the first,

$$.66u(w+2400) - .66u(w) \quad > \quad .66u(w+2400) - .66u(w)$$

- So choice isn't consistent with expected utility maximization, and seems to overvalue the "sure thing")

People overvalue small probabilities (skip?)

- For most people in their sample,

$$.45 \circ \$6000 \quad \prec \quad .90 \circ \$3000$$

  but

$$.01 \circ \$6000 \quad \succ \quad .02 \circ \$3000$$

- (They interpret this as **people overvaluing small probabilities**, and argue this could explain popularity of gambling as well as insurance)

People evaluate gains and losses differently (skip?)

- While they find people are typically risk-averse when it comes to gains,

  they are often **risk-loving with potential losses**

- 92% of respondents said

$$.80 \circ (-\$4000) \quad \succ \quad 1 \circ (-\$3000)$$

  even though it has a lower expected value

- Similarly, 92% said

$$.45 \circ (-\$6000) \quad \succ \quad .90 \circ (-\$3000)$$

- So while utility for money seems to be concave above w, it seems to be convex below 0!

- Kahneman and Tversky propose a different model – basically, using utility functions of different shapes above and below a consumer's starting point, and multiplying each utility term by a function of the probability, not the probability itself – to try to capture the patterns they found

5.3 Discomfort over unknown uncertainty – the Ellsberg paradox

- One other fun example, introduced by Daniel Ellsberg
- a large container contains 300 balls; 100 are red, and the rest are some combination of white and blue, but you don't know exactly how many white and how many blue
- I'm going to pick one ball out of the urn
- First, suppose you're going to get \$100 if you guess the color of the ball correctly; I'll let you guess either red or white
- Most people strictly prefer red
- Second, suppose you get \$100 if you guess a color *other than* the color of the ball
- Again, you can guess red or white
- Again, most people strictly prefer red
- The first suggests people believe there are, on average, fewer white balls than red; the second suggests believe there are more white balls
- In fact, people instead seem to prefer gambles where they know the exact probability, rather than bets where they don't

6 One last point – rational choice

- So yes, particularly when it comes to choice under uncertainty –

  but also in other areas of consumer theory –

  there's a signifcant literature on how real behavior (either experimental or observed) conflicts with our model

- Nobel prize winner Bob Aumann has a recent paper[6]

  in which he argues that we optimize according to evolved decision heuristics –

  he argues these are very close to rational in "real" settings,

  but fail in contrived lab settings that are designed to trip them up

- But that said, in a real sense, we don't "need" people to *always* act rational,

  to believe there's value in understanding a model of rational choice

- David Friedman – son of Milton Friedman,

  an excellent writer and serious libertarian economist – put it well:[7]

  "One summer, a colleague asked me why I had not bought a parking permit.

  I replied that not having a convenient place to park made me more likely to ride my bike.

  He accused me of inconsistency. As a believer in rationality, I should be able to make the correct choice between sloth and exercise without first rigging the game.

  My response was that rationality is an assumption I make about other people.

  I know myself well enough to allow for the consequences of my own irrationality.

  But for the vast mass of my fellow humans, about whom I know very little,

  rationality is the best predictive assumption available."

[6] Aumann (2019), "A Synthesis of Behavioural and Mainstream Economics," *Nature Human Behaviour* 3
[7] David Friedman (1996), *Hidden Order: The Economics of Everyday Life*, pp. 4-5

- For sure, in specific contexts,
  people may behave in seemingly (and even systematically) "irrational" ways;
  and if we know enough about a specific context,
  we might have another model that offers more accurate and precise predictions
- But rational choice gives us a very simple, parsimonious model
  that is "pretty good" over a very wide range of domains –
  and is at very least a good starting point to learn well before building out from

- and that's why we've spent the last 7 weeks thinking about the question I posed day one –
  "everyone optimizes; so what?"
- we've seen the implications when price-taking firms maximize their profits,
  and when price-taking consumers maximize utility given a budget constraint;
  seen what behavior is predicted by these models,
  and therefore what observations would be consistent or inconsistent with them
- and we've at least begun to explore rational choice with uncertainty
- going forward, you all will be looking at similar questions – everyone optimizes, so what –
  when we add in additional effects:
  market-clearing – prices are affected by aggregate supply and demand –
  or strategic concerns – my choices affect other peoples' payoffs and vice versa –
  or information – peoples' behavior reveal information that may be relevant to me
- But the simple case is a key starting point to build out from
- I've enjoyed teaching you, and I wish you luck going forward
- Thanks for being here!