Psychology and Economics

14.13 Lectures 15 and 16: Utility from beliefs; learning

Frank Schilbach

MIT

April 8 and 13, 2020

Plan for next few lectures

- April 8: Utility from beliefs
- April 13: (Non-)Bayesian learning
- April 15: Projection and attribution bias

Why might people miss information and fail to learn?

(1) Attention is limited.

- Too much info in the world to attend to everything.
- Inattention to taxes (Chetty et al., 2009)

(2) People might have wrong theories of the world.

- Rational attention theories: people don't miss important info.
- But might neglect any info if theory says it doesn't matter.
- Learning by noticing (Hanna et al., 2014)

(3) People derive utility from (wrong) beliefs.

- Anticipatory utility
- Ego utility

(4) People might (simply) be bad at (Bayesian) learning.

Utility from beliefs

- Economists typically define utility functions over *outcomes* such as money, consumption, or health.
  - Another source of utility: beliefs about outcomes

- Utility from beliefs can be a very powerful source of utility.
  - Example: high-profile public speech
  - Utility of experience itself might be dwarfed by stress and anticipation derived beforehand from anticipating it (plus the utility from memories).

- Utility from beliefs can even affect physical outcomes.
  - Placebo effect: a treatment can help a patient merely because she believes it will.
  - In fact, it's hard to imagine a source of utility that's not influenced to some extent by beliefs.

Anticipatory utility

- Many emotions are intimately related to what a person thinks about the future: hope, fear, anxiety, savoring, etc.
- Anticipatory utility
  - Utility derived now from anticipating the future
  - Prime example of utility from beliefs
- How might anticipatory utility affect behavior? Two classes of implications:
  (1) Choice of timing of outcomes
  (2) Beliefs and information acquisition

Some motivating evidence of anticipatory utility

- Not so much to establish the existence of such utility, but to start thinking about the patterns of behavior generated by it
- Loewenstein (1987) asked undergrads about (hypothetical) willingness to pay *now* to obtain/avoid certain experiences.
  - WTP as a function of the amount of time until the experience occurs
  - All values normalized relative to students' WTP to obtain the experience right away
- Both pleasant and unpleasant experiences
  - Receiving monetary gains or avoiding monetary losses
  - Obtaining a kiss from a movie star of the student's choice
  - Avoiding a non-lethal but very painful electric shock

When would you like to kiss the movie star?

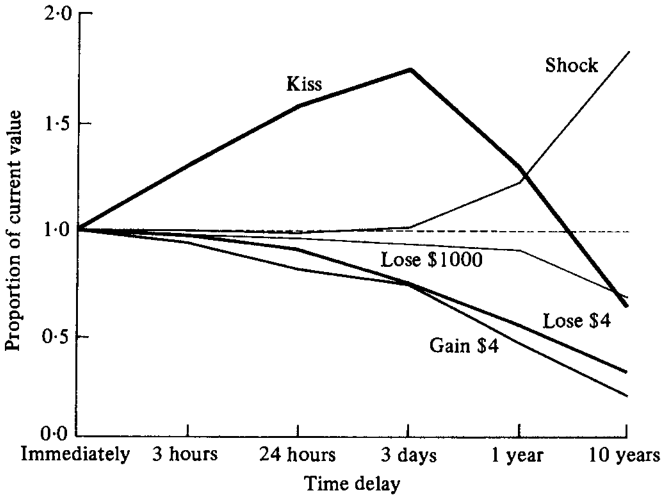

Results

- Subjects prefer:
  (a) a kiss from a movie star in three days rather than now, and
  (b) to have a shock now rather than in one year or ten years.
- This contradicts discounting of any kind!
- Natural explanation:
  - Look forward to the kiss, so delay to enjoy anticipation
  - Unpleasant to anticipate the shock, so get it over with quickly

Why did Loewenstein choose the kiss example?

- Experience with high degree of savorability
- Tries to rule out alternative explanations based on preparation.
- What if he used a dinner at a fancy French restaurant instead?
  - Delay perfectly reasonable even without anticipatory utility
  - You would want to make time, get a date, etc. for the dinner.
- Is preparation problematic for the electric shock example too?
  - No, because that would induce subjects to prefer it later, which is not what Loewenstein finds.

Another example: cleaning hamster cages

- Loewenstein (1987) asked subjects how much they'd have to be paid to clean 100 hamster cages.
- What kind of answer pattern do you expect to see?
- Answer resemble electric shock example!
    - If payment is now for cleaning to be performed next week: \$30.
    - If payment is now for cleaning to be performed in a year: \$37.

Discounting vs. anticipation

- Recall our way of thinking about intertemporal choice:
  (1) What determines a person's instantaneous utility at each moment in time?
  (2) How does she integrate/aggregate those utilities (across time)?
  (3) What does she predict about her future utility and behavior?

- The discounting issues we covered are about point (2)

- The sophistication versus naivete issues are about point (3).

- Anticipatory utility is about point (1), i.e. what enters instantaneous utility.

Interaction of anticipation and discounting

- Anticipatory utility stronger for events closer in time
- For pleasant savorable experiences:
  - Some delay is optimal to have a few periods of anticipation.
  - But not too much for discounting reasons and so as not to weaken anticipatory utility.
- For unpleasant fearful experiences:
  - Either do it immediately to eliminate periods of anticipation, OR:
  - Put it off as much as possible for discounting reasons and to weaken anticipation.

Information gathering and beliefs

- So far: implications of anticipatory utility for *timing* of consumption.
- Perhaps the most important implications of anticipation come from their effect on people's *information gathering* and *beliefs*.
- Would an individual in a non-strategic setting with no anticipatory utility ever strictly prefer to refuse free info?
  - No! The information could be useful for making a decision.
  - Even if not, she can just ignore it (free disposal).
- But a decision-maker with anticipatory feelings can't ignore information, because information affects her emotions.

Information avoidance: Dr. House & Thirteen

(Not) learning about Huntington's Disease

- What is Huntington's Disease?
- Huntington's Disease (HD): degenerative neurological disorder
- Onset around age 40; life expectancy around 60; no cure
- People with parent with HD have 50% chance of developing HD themselves.
- Since 1990s genetic blood test available:
  - Can provide at-risk individuals with certainty about whether they will develop HD.
  - Lab tests cost about $200 to $300 (plus consulting and other costs).
- Testing covered by insurance but most tests are paid out of pocket

Why could this test be valuable even if there is no cure?

- Childbearing
- Marriage
- Retirement
- Education
- Participation in clinical research
- ...

Results from Oster et al. (2013)

- Observe sample of previously untested at-risk individuals over course of ten years
- Low rates of genetic testing: fewer than 10% of individuals pursue predictive testing during the study (Figure 1).
  - Many individuals get tested to *confirm* their disease.
  - Similar evidence in studies on HIV or breast cancer testing
- Untested individuals are (often extremely) overconfident about not getting the disease (Figure 4).

Who gets tested?

812 THE AMERICAN ECONOMIC REVIEW APRIL 2013

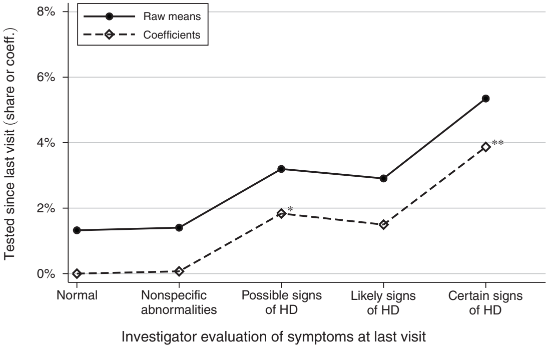

Figure 1. Testing Behavior and Investigator Evaluation of Risk

Notes: This figure shows the relationship between testing and investigator evaluation of risk. The x-axis shows the investigator evaluation of HD status at the last visit. "Possible" signs of HD indicate an estimated 50–89 percent probability of having HD; "Likely" signs indicate 90–98 percent chance; "Certain" signs indicate ≥99 percent chance. The data is limited to individuals who were not tested at the time the evaluation was done (i.e., those at risk of

- Very low overall fraction of people get tested.
- For people with certain signs of HD, objective knowledge from tests is low (since there is no cure).
- Perceived probabilities of having HD might be lower even for people with certain symptoms (next).

Perceived probability of HD as a function of motor scores

VOL. 103 NO. 2 OSTER ET AL.: OPTIMAL EXPECTATIONS AND HUNTINGTON DISEASE 815

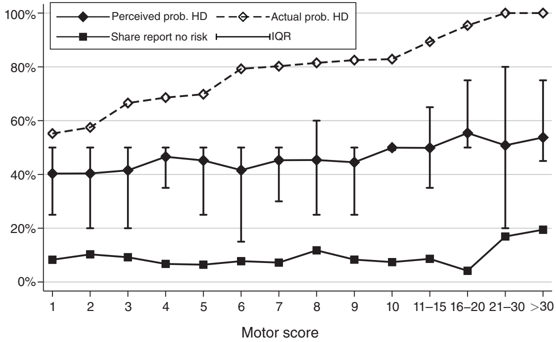

FIGURE 4. PERCEIVED AND ACTUAL RISK OF HD, BY MOTOR SCORE

Notes: This figure shows actual risk of HD by motor score (the dashed line) and perceived risk of HD (upper solid line) by motor score. The bottom line displays the share of individuals at each motor score reporting no risk of HD. The IQR (twenty-fifth to seventy-fifth percentile) for perceived chance is also displayed. Actual risk is based on accurate Bayesian updating based on motor score combined with information on the chance of noticing HD at each symptom

- Untested individuals are often vastly over-optimistic.
- Significant share persist in reporting no chance of having HD.

Do individuals adjust their behavior?

OL. 103 NO. 2 OSTER ET AL.: OPTIMAL EXPECTATIONS AND HUNTINGTON DISEASE 817

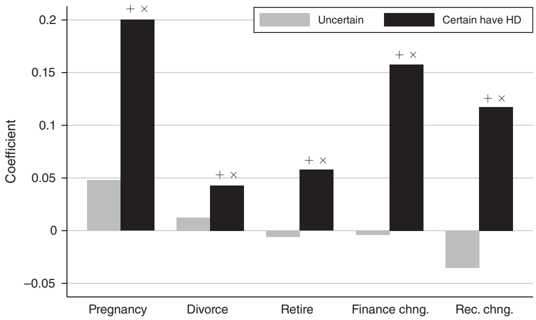

FIGURE 5. BEHAVIOR CHOICE RELATIVE TO INDIVIDUALS WITHOUT HD EXPANSION

Notes: This figure shows differences in behavior relative to tested individuals who report they do not carry the HD mutation. Uncertain individuals are those who are untested and report intermediate probabilities. Those who carry the mutation, know they carry it either through testing or through early symptoms. Regression adjusts for age, education, and gender. The pregnancy regression restricts to people under 40 and controls for existing children and pregnancy last year. + indicates significantly different from those without the HD mutation at 5 percent level;

- Figures shows coefficients relative to those tested negative for HD.
- Those who report certainty about having HD adjust their behavior significantly (see also Table 2 in paper).
- No adjustment for those who report intermediate probabilities (even as high as 90%)

How should we think about these beliefs and behaviors?

- Many people don't get tested despite arguably good reasons to do so.
- People are often over-optimistic about their probability of not having HD, despite often fairly clear signs that they have HD.
- Such over-optimism translates into behavior, i.e. people react less to the signs of likely having HD than they arguably should.
- What models can explain such behaviors?

A simple model

- Two periods, and relevant outcomes occur in period 2.
  - In period 2, the decision-maker will be HD-negative with probability p and HD-positive with probability $1 - p$.
  - Instantaneous utilities are $u(-)$ and $u(+)$, respectively.

- Consider perspective of period 1. Assume:
  (i) No discounting
  (ii) Nothing can be done about the person's condition.

- Expected-utility theory would say that expected utility is:

$$p \cdot u(-) + (1 - p) \cdot u(+).$$

Adding anticipatory utility

- Can only depend on the person's *beliefs* about whether she will have HD, not on what *actually* happens.

- Two extreme assumptions about the formation of beliefs:
  (1) Beliefs are correct.
  (2) Decision-maker can choose (manipulate) her beliefs.

- Start with (1), i.e. the assumption that beliefs are correct.
  - Let anticipatory utility in period 1 be $f(p)$.
  - Recall that p is the probability of being negative. Assume:
    - $u(-) > u(+)$ [obvious]
    - $f(\cdot)$ is increasing [not so obvious].

Does the person want to know her HD status?

- If she doesn't find out her HD status, her utility is

$$f(p) + pu(-) + (1-p)u(+). \tag{1}$$

- If she finds out her HD status, she may find that she is either:
  - HD-negative (with prob p), so her utility is $f(1) + u(-)$; or
  - HD-positive (with prob $1-p$), so her utility is $f(0) + u(+)$.

- Hence, her expected utility is

$$pf(1) + (1-p)f(0) + pu(-) + (1-p)u(+). \tag{2}$$

When does she reject or seek information?

- She will seek information if the expected value from seeking info is higher than the expected value from not seeking info:

$$pf(1) + (1-p)f(0) + pu(-) + (1-p)u(+) > f(p) + pu(-) + (1-p)u(+)$$
$$\Rightarrow pf(1) + (1-p)f(0) > f(p)$$

- Since there is no action (cure), the terms involving $u(-)$ and $u(+)$ are irrelevant for the choice of whether to seek information.

- What determines whether the person gets tested?

When does she reject or seek information?

- She is "information-averse" if

$$pf(1) + (1-p)f(0) < f(p). \tag{3}$$

  - True if $f(\cdot)$ is concave, i.e. steeper for lower values
  - The person *really* dislikes any suspicion of bad news (e.g. due to anxiety), but there is not much added value of certainty.

- She is "information-loving" if

$$pf(1) + (1-p)f(0) > f(p). \tag{4}$$

  - True if $f(\cdot)$ is convex, i.e. steeper for higher values.
  - The person *really* likes certainty; some suspicion of bad news is not so painful.

Belief manipulation

- Now consider the possibility that the person can manipulate her beliefs to make herself feel better.

- In the above framework, would she *want* to hold correct beliefs?
  - No. Why?
  - If she could believe that she is HD-negative for sure, she would get higher utility whatever actually happens later:

$$f(1) + pu(-) + (1-p)u(+) > f(p) + pu(-) + (1-p)u(+) \quad \text{for any } p < 1.$$

- Why might she not want to choose $p = 1$ anyway?

Incorrect beliefs can lead to mistaken decisions.

- Overly positive beliefs are an economically important implication of utility from beliefs.
- Optimal expectations trade off anticipatory utility vs. value of knowledge for making correct choices
  - Person *wants to* believe she is healthy (it makes her feel better), so she convinces herself that that's the case.
  - But over-optimism distorts decision-making (e.g. health behavior). It might also lead to disappointment.
- In general, for a decision-maker with anticipatory utility, at least *some* overoptimism leads to higher utility than realism.

Other evidence of overly optimistic beliefs

- Weinstein (1980) asked students to make judgments of their own and other students' chances for a number of outcomes.
- Two measures:
  (1) **Comparative judgment**: how much more or less likely the average student thinks the event will happen to them relative to the average student.
  (2) **Optimistic/pessimistic ratio**: number of students who think their chances are better than the average classmate's, divided by number who think their chances are worse.

Results from Weinstein (1980) study: clear evidence of over-optimism

Event	**Comparative judgment (%)**	**Optimistic/pessimistic**
Like post-graduation job	50.2***	5.93***
Owning house	44.3***	6.22***
Starting salary > $10,000	41.5***	4.17***
Traveling to Europe	35.3***	2.25***
Living past 80	12.5**	2.00**
Having a drinking problem	-58.3***	7.23***
Attempting suicide	-55.9***	8.56***
Divorce soon after marriage	-48.7***	9.50***
Heart attack before age 40	-38.4***	5.11***
Having gum problems	-12.4**	1.39

Many other forms of positively biased beliefs

- Other examples:
  (1) Couples believe there is a small chance their marriage will end.
  (2) Small-business owners think their business is far more likely to succeed than a typical similar business.
  (3) Smokers understand the health risks of smoking, but don't believe those risks apply specifically to them.

- Beyond future prospects, people tend to have overly positive views about their abilities and traits.
  - 90% of drivers think they are better than the average driver.
  - 94% of professors at the University of Nebraska think they are better teachers than the average professor at their university.

In contrast, MIT students appear under-confident.

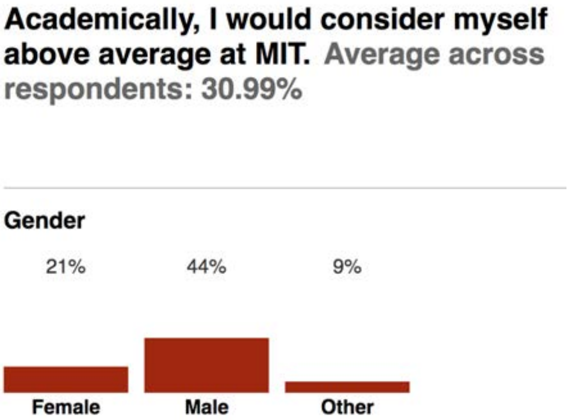

- Why?
- Worries about disappointment?
- Biased beliefs about the average?
- Other reasons?

Female MIT students are particularly under-confident.

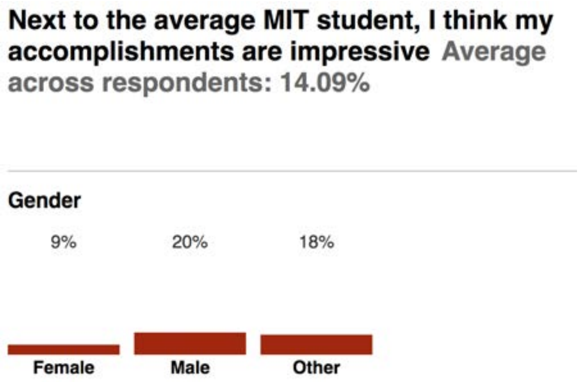

- Men tend to be more overconfident than women more generally.
- See, for instance, Barber and Odean (2001) below.

Ego utility

- Inflated beliefs might be explained by anticipatory utility:
  - Since higher ability means better future prospects, people want to convince themselves that they have high ability.
  - Unclear whether traits are closely related to future prospects.
  - Likely that people have another type of utility from beliefs
- Ego utility
  - Utility derived from beliefs about one's traits and abilities.
  - Then, as with anticipatory utility, people may want to hold rosy views of their abilities because it makes them feel better.

Aside: overoptimism and mental health

- Taylor and Brown (1988): positive illusions promote psychological wellbeing
- Argument: overconfidence is vital for mental health.
  - Happiness
  - Ability to care for others
  - Capacity for creative, productive work
  - Management of negative feedback
  - ...
- Corollary: realistic expectations might be detrimental to mental health!
- Will return to this issue in lecture on happiness and mental health

Some factors that affect the extent of positive biases

- People tend to have *greater* biases with respect to prospects and traits that are personally important to them.
- Available or imminent objective info tends to *decrease* biases.
- If feedback about the prospect is more ambiguous and subject to interpretation, biases tend to be *greater*.
- If people feel like they have control over the outcome, biases tend to be *greater*.
- Expertise sometimes *increases* biases, but not for experts who get very good feedback (e.g. meteorologists).

Do people act on biased beliefs?

- Possible concern: above evidence is all about self-reports
  - Are beliefs biased when people are incentivized to give correct answers?
  - Do biased beliefs translate into suboptimal choices?

- Lab experiment by Eil and Rao (2010)
  - Feedback about IQ test score and physical attractiveness
  - Elicit prior belief about rank between 1 and 10
  - Get bilateral comparison with one other participant in group
  - Then elicit posterior and willingness to pay for true rank

- Asymmetric processing of objective information about self
  - Favorable news: subjects are roughly Bayesian (slightly over-optimistic)
  - Unfavorable news: discount signals; noisy posterior beliefs

Further evidence of real-world consequences of biased beliefs

- Moebius et al. (2014): more asymmetric belief updating
  - Evidence similar to Eil and Rao (2010)
- Zimmermann (2019): motivated memory
  - Memory can play an important role in the formation of motivated beliefs.
- Barber and Odean (2011): overconfident small-scale investors
  - Men trade 45% more than women.
  - They lose money, plausibly due to overconfidence.
- Malmendier & Tate (2006): managerial hubris
  - Clever way of identifying overconfident managers: (long)holding stock options
  - Overconfident managers engage their businesses in more mergers (which on average destroy value)

Why might people miss information and fail to learn?

- Attention is limited.
  - Too much info in the world to attend to everything.
  - Inattention to taxes (Chetty et al., 2009)
- People might have wrong theories of the world.
  - Rational attention theories: people don't miss important info.
  - But might neglect any info if theory says it doesn't matter.
  - Learning by noticing (Hanna et al., 2014)
- People derive utility from (wrong) beliefs.
  - Anticipatory utility
  - Ego utility
- People might (simply) be bad at (Bayesian) learning.

Economic decision-making under uncertainty

- Almost all meaningful decisions are undertaken under at least some *uncertainty*.
  - We cannot know for sure what will happen.
  - However, it is possible to know something about the *likelihood* of relevant events.

- Examples:
  - Student decides which topics in the course to focus on
  - Basketball coach decides whether to leave tired player in game
  - Many medical, managerial, educational, and career decisions.

- How do individuals make probability judgments of this kind?
  - So far: risk preferences for given probabilities (e.g. **?**)
  - Now: how do people learn about such probabilities?
  - Pioneering work by **?**

What we get right

- In terms of broad directions, people's probability judgments are very good.
- Example: imagine you're deciding whether to see the new James Cameron movie.
    - You read an online review.
    - You heard various opinions from your friends.
    - You look up the Rotten Tomatoes rating.
- How will you form an estimate of the likelihood that you will enjoy the movie?

How likely is that you will like the movie?

- Likely estimation procedure:
  (1) Consider your experience with James Cameron movies.
  (2) This determines a baseline likelihood ("base rate") of the new movie being good.
  (3) Determine whether and how the tastes of the online review, your friends and Rotten Tomatoes are related to your own.
  (4) You can use these other opinions to adjust your assessed probability of the new movie being for you.

- This procedure is helpful regardless of whether others' tastes are positively *or* negatively related to your own.

- Completely reasonable and mathematically well-founded procedure

Intuitive likelihood estimates

- Most people understand basics of forming likelihood estimates.
- They know the direction in which features of the baseline likelihood and of available information should affect estimates.
    - If you liked more James Cameron movies in the past, then, holding the opinions of the critic and your friends constant, it's more likely that you'll like the new one.
    - If your friends' opinions are typically close to your own, you should adjust your estimate toward their opinions more than if your opinions are almost unrelated.

What we get wrong: putting together the pieces

- While in broad terms intuitive likelihood judgments are in the right direction, they're not very accurate.
  - Understanding the direction different pieces of information should pull our estimates is easy.
  - Knowing *how far* to adjust estimates is very difficult.
- Applying Bayes' Rule is cognitively extremely demanding, so that in most situations people don't get it completely right.

Bayes' Rule through an example

- Suppose you have a coin that you start off thinking is fair with probability $2/3$.
- The coin is biased towards heads with probability $1/3$, in which case heads comes up 75% of the time.
- You flip the coin and it comes up H. What's the new probability that it's fair?
  - Clearly, it should be less than $2/3$.
  - But by how much? That's what Bayes' rule tells us.

Graphical illustration

- Suppose there are two types of coins:
  (1) Fair coin: an urn with an equal number of H and T balls.
  (2) Unfair coin: an urn in which 75% of the balls are H.
- You don't know whether the coin is fair, i.e. you don't know which urn you're drawing from.
- One way to think about this problem: imagine you're drawing from an urn containing both urns.

Bayes' Rule

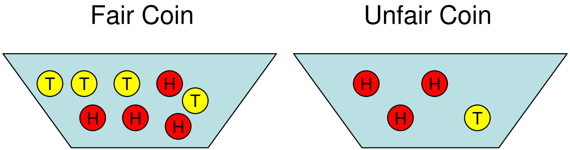

- Initial probability of a fair ball is 2/3, i.e. the fair urn has twice as many balls.
- Suppose we have drawn a ball randomly and it's H. What's the probability that it came from the fair urn? It's 4/7.
- If we draw a ball randomly and it's T. The probability that it came from the fair urn is 4/5.

Bayes' Rule

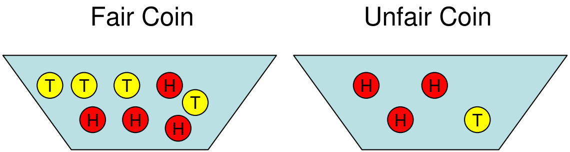

- The probability of a hypothesis h being true, *conditional on observing* information i, is

$$
\begin{aligned}
Prob[h|i] &= \frac{\text{number of ways } i \text{ and } h \text{ can both be true}}{\text{number of ways } i \text{ can be true overall}} \\
&= \frac{\text{probability that } i \text{ and } h \text{ are both true}}{\text{probability that } i \text{ is true}}.
\end{aligned}
$$

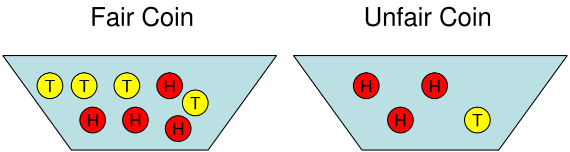

- Simple way to remember Bayes' Rule:

$$P(h \cap i) = P(i|h) \cdot P(h) = P(h|i) \cdot P(i)$$
$$\Rightarrow P(h|i) = \frac{P(i|h) \cdot P(h)}{P(i)} = \frac{P(h \cap i)}{P(i)}$$

- Here:

$$\frac{P(h \cap i)}{P(i)} = \frac{\frac{1}{2} \cdot \frac{2}{3}}{\frac{1}{2} \cdot \frac{2}{3} + \frac{3}{4} \cdot \frac{1}{3}} = \frac{\frac{4}{12}}{\frac{7}{12}} = \frac{4}{7}$$

Biases in probability judgments: an overview

- Even this simple example of Bayes' Rule is somewhat difficult to follow.
- *Applying* the rule in a *new* situation with *multiple* pieces of different information is far more difficult.
- Most people don't know the precise rule, and even if they do they can't or don't want to think so hard.
- Instead, they use intuitive shortcuts to make judgments of likelihood.

Shortcuts lead to systematic mistakes.

- People attempt to form likelihood estimates by focusing on an aspect of the situation at hand that seems most relevant.
- Because they ignore other relevant aspects of the situation, their likelihood estimates will typically be incorrect.
- Because different economic situations lead individuals to focus on systematically different aspects, intuitive judgments lead to *systematic* mistakes in judgments of probability.
- What do people focus on, and what systematic mistakes does this lead to?

Many real-world situations involve sequences of random outcomes

- Decision-makers often observe sequence of related outcomes before making choice between uncertain options.
- An investor might observe past performance of mutual funds before deciding which one to invest in.
- A patient or doctor might observe the outcome of prior surgeries before deciding whether to undertake the surgery.
- A coach can observe the recent performance of a basketball player before deciding whether to put the player in the game.

Pattern I: The Gambler's Fallacy

- False belief that in a sequence of independent draws from a distribution, an outcome that hasn't occurred for a while is more likely to come up on the next draw.
- "If red came up in roulette four times in a row, black is due."
- Almost folk knowledge

Careful experiment by Gold and Hester (1987)

- Subjects were told that a coin with one black and one red side would be flipped 25 times.

- The experimenter actually reported a pre-determined sequence of flips: 17 mixed, then 1 black, and then 4 red.

- On the 23rd flip, participants were given a choice between
  (i) 70 points for sure;
  (ii) 100 points if the next flip was "their color."

- Half of the subjects' color was red and half's was black.

- Propensity to take the 70 points reveals beliefs about the odds that the next flip would be their color.

Gambler's Fallacy at work

- Clear beliefs that next flip was more likely to be black
  - 24 of 29 "red" subjects chose to take the sure thing
  - 8 of 30 "black" subjects did.
- An interesting variant:
  - For some participants, 23rd coin flip delayed by 24 mins
- Weaker evidence of Gambler's Fallacy:
  - 18 of 32 "red" subjects chose the sure thing
  - 13 of 29 "black" subjects chose the sure thing.
  - Subjects seem to believe that letting a coin "rest" makes it more likely that a streak would continue!

Pattern II: The Hot-Hand Fallacy

- Basketball fans/players/coaches believe there is *systematic* day-to-day variation in players' shooting performance.

- Performance of a player may sometimes be predictably better than expected on the basis of the player's overall record.
  - Example: "She is on fire today."; or "He is a streak shooter."

- Presumably, "on fire today" means that he's more likely to hit his shots today than on other days.
  - Made shots should cluster together
  - Should be more likely to hit next shot if hit recent shots
  - Lots of work on this issue, starting with Gilovich, Vallone & Tversky (1985)
  - Read more about this HERE.

A common explanation?

- In a way, the Hot-Hand Fallacy is exactly the opposite of the Gambler's Fallacy.
  - The Gambler's Fallacy is the belief that the next outcome is likely to be *different* from the previous ones.
  - The Hot-Hand Fallacy is the belief that the next outcome is likely to be *similar* to the previous ones.
- In fact, both can be thought of as a consequence of the "Belief in the Law of Small Numbers."
- Law of Large Numbers
  - Mathematical fact: in large sample of independent draws from a distribution, the proportions of different outcomes closely reflect their underlying probabilities.

Belief in the Law of Small Numbers (BLSN)

- False belief that in *small* samples, the proportions of different outcomes should reflect their underlying probabilities.
- While false, this belief is a natural consequence of focusing too much on a central aspect of the situation at hand.
- Belief that the proportions of different outcomes in a sequence of random draws are related to their underlying probabilities.

BLSN can explain both gambler's fallacy and hot-hand fallacy

- When a person *knows* (or is confident about) the distribution, BLSN tends to generate the Gambler's Fallacy.
  - Four reds in a row in roulette doesn't look like the expected sequence of half black and half red.
  - Four reds, one black is closer, so she predicts black on next flip.
- When a person *doesn't know* the distribution, BLSN can lead to the Hot-Hand Fallacy.
  - A streak of made shots looks like a sequence one would expect from a player on fire, i.e. a high probability of making shots.

Pattern III: Base-Rate Neglect

- A question from the survey in the first lecture:
  - *One in a hundred people have HIV, and we have a test for HIV that is 99% accurate. If a person tested positive, what's the probability that she has HIV?*
  - True answer is 50%, but many of you answered 99%.
  - You probably forgot to take into account how few people in the population have HIV.
- **Base-Rate Neglect:** Ignoring or underweighting the base rate of an event in judging its likelihood when presented with information about it.
  - Natural consequence of focusing too much on a central aspect of situation at hand.
  - When presented with information about a hypothesis, people focus on the implications of the information. As a consequence, they ignore the base rate.

Remember Bayes' Rule?

- Notation P=HIV-positive, N=HIV-negative, and p=tested positive

$$\begin{aligned}
\Pr[P|p] &= \frac{\text{probability that } P \text{ and } p \text{ are both true}}{\text{probability that } p \text{ is true}} \\
&= \frac{\Pr(P\&p)}{\Pr(p)} = \frac{\Pr(P\&p)}{\Pr(P\&p)+\Pr(N\&p)} \\
&= \frac{(.01)(.99)}{(.01)(.99)+(.99)(.01)} = 0.5
\end{aligned}$$

Fake news: two-dimensional learning problems are inherently hard.

AMERICA

"Students Have Dismaying Inability to Tell Fake News from Real, Study Finds." by Camila Domonoske. NPR. Nov. 23, 2016.

- Need to learn not just from sources wit given accuracy but also about the accuracy of those sources
- Recent work by David Rand and coauthors considers strategies to reduce belief in fake news.

Misunderstanding probabilities: summary

- Using Bayes' Rule is difficult.
- Systematic deviations from Bayesian updating:
  (1) Gambler's Fallacy
  (2) Hot Hand Fallacy
  (3) Base-Rate Neglect
- (1) and (2) can be explained by the belief in the "Law of Small Numbers."
- Some real-world implications:
  - Decision-making by judged, umpires, loan officers (Chen, Moskowitz & Shue 2016)
  - Even highly accurate tests for Covid-19 antibodies are not very informative if relatively low fraction of people infected.

Heuristics and Biases (Kahneman & Tversky)

- Way to think about biases in probability judgments. Starting point:
  - Applying the laws of probability and statistics is often impossibly hard.
  - So people use their quick and intuitive judgments to make likelihood estimates.
- Seminal work by **?** and lots of subsequent work
  - Seek to understand statistical thinking that captures how these intuitions work.
  - Read "Thinking, Fast and Slow" if you'd like to learn more.
- Judgment Heuristic: An informal algorithm that generates an approximate answer to a problem quickly (hard to model)

Heuristics have a good and a bad side

- Heuristics speed up/make possible cognition.
- But because they are shortcuts, they occasionally produce incorrect answers, or "biases."
- These are unintended side effects of adaptive processes.
- It's useful to study heuristics and biases together, much like studying vision and optical illusions together is useful.

The Representativeness Heuristic

- *Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.*

- Rank the following statements from most to least probable:
  (1) Linda is a teacher in elementary school.
  (2) Linda works in a bookstore and takes Yoga classes.
  (3) Linda is active in the feminist movement.
  (4) Linda is a psychiatric social worker.
  (5) Linda is a member of the League of Women Voters.
  (6) Linda is a bank teller.
  (7) Linda is an insurance salesperson.
  (8) Linda is a bank teller and is active in the feminist movement.

The Representativeness Heuristic

- *Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.*

- Rank the following statements from most to least probable:
  (1) Linda is a teacher in elementary school.
  (2) Linda works in a bookstore and takes Yoga classes.
  (3) Linda is active in the feminist movement.
  (4) Linda is a psychiatric social worker.
  (5) Linda is a member of the League of Women Voters.
  (6) Linda is a bank teller.
  (7) Linda is an insurance salesperson.
  (8) Linda is a bank teller and is active in the feminist movement.

Getting confused about Linda

- Subjects rank it as more likely that Linda is both a bank teller and a feminist than that she's a bank teller!
  - Violation of conjunction rule, a basic law of probability
  - The probability that two events both happen is never higher than the probability that just one of the two events happens.

- Potential concern:
  - Subjects might interpret "Linda is a bank teller" as "Linda is a bank teller and is *not* active in the feminist movement."

- Can run the following experiment:
  - Group A sees list without the conjunctive event, item #8.
  - Group B sees list without the non-conjunctive event, item #6.
  - Still, Group B ranks #8 higher than Group A ranks #6.

Representativeness Heuristic

- People use similarity or representativeness as a proxy for probabilistic thinking.
    - Based on the available information, they form a mental image of what Linda is like.
    - When asked about how likely Linda is a school teacher, they ask: how similar is my picture of Linda to a school teacher?
    - They then turn this similarity judgment into a probability judgment.
- Very reasonable heuristic:
    - Easy to form a mental image of an instance, and to judge similarity to a population
    - Often things are more likely to come from populations that are similar to them.
    - But the heuristic can also lead to mistakes.
    - Similarity is sometimes a poor predictor of true probability.

Availability Heuristic

- People assess the probability of an event by the ease with which instances or occurrences can be brought to mind.

- Example: "Are there more suicides or more homicides in the US each year?"
    - You might think about instances of suicides and homicides you can recall, and judge the frequency of each based on how many instances you can recall.
    - This leads most people to think that murders are more common.

- Very sensible heuristic: more often than not, it's easier to recall things that are more common or probable.

Familiarity

- But the heuristic also generates biases because a class whose instances are more easily retrieved will appear more numerous.
- Recall is aided by salience of memories due to rehearsal:
  - Example: perceived vs. actual death rates (for different causes)
  - Memories that are brought to mind many times come to mind more easily again.

Cause	Estimate	Truth
Tornado	564	90
Fireworks	160	6
Asthma	506	1886
Drowning	1684	7380

  - Due to media exposure, we rehearse some memories/issues much more than others.
  - Familiar/personally important events are easier to recall.

Anchoring and Adjustment Heuristic

- People often try to answer a question by starting at some first-pass guess based on memory or the environment.

- Then adjust that guess until they are satisfied with the answer.
  - Very sensible strategy, but people's judgment is colored too much by their original guess (the "anchor").
  - Classic demonstration:
    - Subjects shown blurred pictures that were gradually brought into sharper focus
    - Different subjects began viewing the pictures at different points in the focusing process.
    - Half as many of those subjects who began their viewing at a severe-blur stage eventually identified the pictures correctly as did those who began viewing at a light-blur stage.

- Might help explain why fake news tend to stick.

Summary

- We studied several reasons why people might miss information and fail to learn
  (1) Attention is limited.
  (2) People might have wrong theories of the world.
  (3) People derive utility from (wrong) beliefs.
  (4) People might (simply) be bad at (Bayesian) learning.

- Important to understand underlying reasons why people are misinformed
  - Vastly different policy implications
  - Making information salient (e.g. food labels) makes sense if we think people miss relevant and important information.
  - But providing information will have little effect if people don't want to learn.
  - If people derive utility from beliefs, correcting those beliefs can make them worse off!
  - Understanding systematic biases based on heuristics can help improve decisions.

What's next?

- Projection and attribution bias
- Read for Wednesday: Loewenstein et al. (2003), sections I, II, and III

References used in this lecture I

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Hanna, Rema, Sendhil Mullainathan, and J. Schwartzstein, "Learning Through Noticing: Theory and Evidence from a Field Experiment," *Quarterly Journal of Economics*, 2014, *129* (3), 1311–1353.

Kahneman, Daniel and Amos Tversky, "Judgment under Uncertainty: Heuristics and Biases," *Science*, 1974, *185* (4157), 1124–1131.

__ and __ , "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 1979, *47*, 263–292.

Loewenstein, George, "Anticipation and the Valuation of Delayed Consumption," *Economic Journal*, 1987, *97* (387), 666–684.

__ , Ted O'Donoghue, and Matthew Rabin, "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics*, 2003, *118* (4), 1209–1248.

Oster, Emily, Ira Shoulson, and E. Ray Dorsey, "Optimal Expectations and Limited Medical Testing: Evidence from Huntington's Disease," *American Economic Review*, 2013, *103* (2), 804–830.

Weinstein, Neil, "Unrealistic Optimism About Future Life Events," *Journal of Personality and Social Psychology*, 1980, *39* (5), 806–820.