

## **Economics 142: Probabilistic judgment**

**Winter 2008**

**Vincent Crawford (with large debts to Colin Camerer and especially Botond Koszegi and Matthew Rabin)**

Most economic decisions are made in dynamic settings with some uncertainty, and so require probabilistic judgment to draw correct inferences from observations and form correct beliefs.

The standard economic assumption has been that people make such decisions optimally, using the laws of probability theory. But just as considering evidence on choice behavior led us to question some of the standard assumptions about preferences, evidence can also lead us to question standard assumptions about probabilistic judgment and the formation of beliefs.

Led by Kahneman and Tversky (1974 *Science*) and others, psychologists and economists have used deviations between intuitive probability judgments and normative principles (“biases”) to suggest general principles of how probabilistic judgment deviates from rationality.

Their approach was inspired by theories of perception, which use optical illusions to suggest principles of vision without implying that everyday visual perception is badly maladaptive.

Although the use of heuristics and the resulting biases can lead to choices that are suboptimal when judged by the standards of mainstream economics, the point of this line of research is not to argue that humans are stupid: Heuristics need not be maladaptive.

## Generalizations from Tversky and Kahneman's "Heuristics and Biases" (*Science* 1974)

Tversky and Kahneman (1974) conceptualize observed departures from perfect rationality in judgment under uncertainty by noting that

... people rely on "heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors."

Tversky and Kahneman also provide an apt metaphor for this approach:

The subjective assessment of probability resembles the subjective assessment of physical quantities such as distance or size. These judgments are all based on data of limited validity, which are processed according to heuristic rules. For example, the apparent distance of an object is determined in part by its clarity. The more sharply the object is seen, the closer it appears to be. This rule has some validity, because in any given scene the more distant objects are seen less sharply than nearer objects. However, the reliance on this rule leads to systematic errors in the estimation of distance. Specifically, distances are often overestimated when visibility is poor because the contours of objects are blurred. On the other hand, distances are often underestimated when visibility is good because the objects are seen sharply. Thus, the reliance on clarity as an indication of distance leads to common biases. Such biases are also found in the intuitive judgment of probability.

Tversky and Kahneman's viewpoint is that there are two interacting systems in cognition:

- The intuitive system uses heuristics that sometimes get things wrong from the point of view of conscious reasoning; but it is fast, automatic, effortless, and difficult to control or modify. It is adaptive because it gets things approximately right when it is important to act quickly.
- Conscious reasoning is a slow but sophisticated process that is very flexible and can be changed and improved by learning; but it can only concentrate on one thing at a time, and it requires effort and control.

Choice is the product of a continual interaction between these two systems, in which conscious reasoning struggles to override intuition, But even when the evidence against intuition is strong, it fights back and sometimes wins.

Errors or biases in judgment are the unintended side effects of generally adaptive processes.

Most probability theory is plainly more in the realm of conscious reasoning than intuition, so intuitive probability judgments are unlikely to be fully rational.

Just as one can learn a lot about vision from optical illusions, one can learn a lot about probabilistic judgment by studying its biases. Because the mainstream alternative is assuming that people have perfect probabilistic judgment, there is considerable value in studying biases.

## **Bayes' Rule (Koszegi)**

The key component of rational probabilistic judgment is Bayes' Rule. To understand Bayes' rule, suppose you have a coin, but you do not know whether it is fair. You start off thinking that it is fair—so that it gives heads 50% of the time—with probability two-thirds, and it is biased toward heads—so that it gives heads 75% of the time—with probability one-third.

Now imagine that you flip the coin and it comes up heads. How would you have to change your beliefs about the probability that the coin is fair? You would clearly have to decrease it, since a head outcome is more likely to come from a biased coin. But by how much?

A fair coin is like an urn in which exactly half the balls say heads, and half say tails: the probability of getting heads on a flip of the coin is the same as the probability of getting a “heads” ball when drawing randomly from the urn.

The biased coin is like an urn in which 75% of the balls say heads, and 25% say tails.

You are drawing a ball from one of these urns, but you do not know which one. So imagine you are drawing a ball from one big urn, with the two small urns combined inside it.

What does it mean that you think the coin is fair with probability two-thirds? It means that the urn in which exactly 50% of the balls say heads and 50% say tails has twice as many balls, in total, as the one in which 75% of the balls say heads and 25% say tails. In other words, when drawing a ball from the big urn, you think it is twice as likely to come from the fair urn as from the unfair urn.

These probabilities can be represented by letting the unfair urn have four balls in total, three heads and one tails, and letting the fair urn have eight balls in total, four heads and four tails.

We can now see how much you should change your beliefs about the probability that the coin is fair if you flip the coin once and it comes up heads. Suppose we draw a ball from the big urn and it says heads. What is the probability that it came from the fair urn? There are four heads balls in the fair urn and three in the unfair one, for a total of seven. So the probability is 4/7.

We have just derived Bayes' rule for this example: More generally, the probability of hypothesis  $h$  being true, conditional on observing information  $i$ , is  $\text{Prob}[h|i] = \text{Prob}[i \text{ and } h] / \text{Prob}[i]$ .

In the coin example, the hypothesis  $h$  is that the coin is fair. The information is that when flipped once, it came up heads. The number of ways hypothesis  $h$  and information  $i$  can both be true is the number of ways to draw a heads ball from the fair urn—that is, 4. The number of ways  $i$  can be true is the number of ways to draw a heads ball from the big urn—that is, 7.

## Base rate neglect

Now considering some experimental evidence on probabilistic judgment, beginning with survey questions 3 and 4 (a and b). Both fall under the heading of a bias called base rate neglect.

Survey question 3: Suppose that one out of a hundred people in the population have HIV. There is a test for HIV that is 99% accurate. This means that if a person has HIV, the test returns a positive result with 99% probability; and if a person does not have HIV, it returns a negative result with 99% probability. If a person's HIV test comes back positive (and you know nothing else about her/him), what is the probability that s/he has HIV?

Most people answer 99%. This is wrong. The correct answer, 50%, can be found in two ways.

- The intuitive argument in the Introduction is one: An HIV-negative person is 99 times less likely to test positive than an HIV-positive person, but there are 99 times more HIV-negative people. These cancel out, so the probability that a person testing positive has HIV is 50%.

- Bayes' Rule is a more systematic way: It says that the probability that a person whose HIV test comes back positive has HIV is the ratio of the probability that {a person's test comes back positive and the person has HIV} to the probability that {a person's test comes back positive}.

The probability that {a person's test comes back positive and the person has HIV} under the stated conditions is  $0.01 \times 0.99$ , because 1% of the population have HIV and the test gives the right answer for 99% of them.

The probability that {a person's test comes back positive} is the sum of two terms: the same  $0.01 \times 0.99$  from the people who do have HIV, plus another  $0.99 \times 0.01$  from the 99% of the population who don't have HIV and the test gives the wrong (positive) answer for 1% of them.

Thus Bayes' Rule gives the probability that a person whose HIV test comes back positive actually has HIV as  $(0.01 \times 0.99) / [(0.01 \times 0.99) + (0.99 \times 0.01)] = \frac{1}{2}$ , like the intuitive argument.

As both Bayes' Rule and the intuitive argument show, the problem with the common answer is that it ignores the base rate ("one out of a hundred people in the population have HIV"). Unless the test is perfectly accurate, even intuition suggests that the base rate is relevant (though not how to combine it with the test result), but people tend to ignore it anyway.

Ignoring the base rate makes people systematically overestimate the probability of rare events—such as a person having HIV—because it ignores their rarity; and it makes them underestimate the probability of common events—such as a person not having HIV.

To see base rate neglect in a different way, consider survey questions 4 (a and b):

4a. Jack's been drawn from a population which is 30% engineers and 70% lawyers. Jack wears a pocket protector. Use your own estimate of the respective probabilities that engineers and lawyers wear pocket protectors to estimate the probability  $p_1$  that Jack is an engineer.

4b. Jack's been drawn from a population which is 30% lawyers and 70% engineers. Jack wears a pocket protector. Use your own estimate of the respective probabilities that lawyers and engineers wear pocket protectors to estimate the probability  $p_2$  that Jack is an engineer.

People's average estimates for  $p_1$  and  $p_2$  are virtually the same.

The right estimates depend on the estimated probabilities that lawyers and engineers wear pocket protectors, which we don't know; but even so we can tell that it's wrong to have the same estimates for  $p_1$  and  $p_2$ :

If  $q$  is your estimated probability that a lawyer wears a pocket protector and  $r$  is your estimated probability that an engineer wears a pocket protector, then, using Bayes' Rule again,  $p_1 = 0.3r/[0.3r+0.7q]$  and  $p_2 = 0.7r/[0.7r+0.3q]$ .

Thus  $p_1/(1-p_1) = 0.3r/0.7q$  and  $p_2/(1-p_2) = 0.7r/0.3q$ , so  $[p_1/(1-p_1)]/[p_2/(1-p_2)] = (0.3/0.7)^2 \approx 18\%$ . ( $q$  and  $r$  cancel out.)

Again, people ignore the base rate and so systematically overestimate the probability of the rare event (engineer dresses well) and underestimate the probability of the common event.



## Representativeness

Consider the following example (Kahneman and Tversky, 1973 *Psychological Review*):

“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.”

Please rank the following statements by their probability, using 1 for the most probably and 8 for the 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.

I now repeat them with typical average probability rankings in parentheses:

1. Linda is a teacher in elementary school. (5.2)
2. Linda works in a bookstore and takes Yoga classes. (3.3)
3. Linda is active in the feminist movement. (2.1)
4. Linda is a psychiatric social worker. (3.1)
5. Linda is a member of the League of Women Voters. (5.4)
6. Linda is a bank teller. (6.2)
7. Linda is an insurance salesperson. (6.4)
8. Linda is a bank teller and is active in the feminist movement. (4.1)

Note that item 8 is ranked more likely than item 6; depending on the subject population (but even extending to professional statisticians), 70%-90% rank them in this order.

This is wrong! (Why?)

Kahneman and Tversky call this the conjunction effect (since the conjunctive event receives a higher probability.)

This same phenomenon shows up in many other forms.

Why does it happen? Kahneman and Tversky and others have argued that it's because decision makers use similarity as a proxy for probability. Based on the available information, they form a mental image of what Linda is like. When asked about the likelihood that Linda is a school teacher, bank teller, feminist, and so on, they ask themselves how similar is my picture of Linda to a typical school teacher, bank teller, or feminist? They then turn this similarity judgment into a probability, with more similarity implying a higher probability.

The similarity between a Linda and a feminist bank teller is greater than the similarity between Linda and a bank teller, so they judge item 8 as more likely than item 6.

But by the laws of probability, the probability that Linda is a feminist bank teller must be less than the probability that she is a bank teller—the conjunction rule.

The problem is that similarity relations do not follow the conjunction rule, a basic law of probability.

Base rate neglect is closely related to representativeness.

## Sample size neglect

Consider the following example (Kahneman and Tversky, 1974 *Science*):

A certain town is served by two hospitals. In larger hospital, 45 babies are born per day. In smaller hospital, 15 babies are born per day. 50% of babies are boys, but the exact percentage varies from day to day. For a period of 1 year, each hospital recorded the days on which more than 60% of the babies born were boys.

Which hospital do you think recorded more such days?

- The large hospital?
- The small hospital?
- About the same (within 5% of each other)

Most people think they are about the same.

This is wrong! (Why?)

Why does it happen? Also closely related to representativeness: Subjects assess the likelihood of a sample result by asking how similar that sample result is to the properties of the population from which the sample was drawn.

## “Law of small numbers”

All families of six children in a city were surveyed. In 72 families the exact order of births of boys and girls was GBGBBG. What is your estimate of the number of families surveyed in which the exact order of births was BGBBBB?

In standard subject pools the median estimate is 30. This is wrong! (Why?)

Why does it happen? Also closely related to representativeness:

- People expect that a sequence of events generated by a random process will reflect the essential characteristics of that process even if the sequence is short. (The law of large numbers says that very large samples drawn from a probability distribution very accurately reflect the probabilities in the distribution. People mistakenly apply the same idea to small samples.)
- So if a coin is fair, subjects expect HHH to be followed by a T (the gambler’s fallacy: the false belief that in a sequence of independent draws from a distribution, an outcome that has not occurred for a while is more likely to come up on the next draw).
- If girls are as likely as boys, subjects expect GGG to be followed by B.
- So BGGBBG is viewed as a much more likely sequence than BBBBBB.
- People expect that the essential characteristics of the process will be represented, not only globally in the entire sequence, but also locally in each of its parts.

## More on sample-size neglect and the “law of small numbers” (Rabin)

Griffin and Tversky (1992) experiment: Students read the following.

“Imagine that you are spinning a coin, and recording how often the coin lands heads and how often the coin lands tails. Unlike tossing, which (on average) yields an equal number of heads and tails, spinning a coin leads to a bias favoring one side or the other because of slight imperfections on the rim of the coin (and an uneven distribution of mass). Now imagine that you know that this bias is  $3/5$ . It tends to land on one side 3 out of 5 times. But you do not know if this bias is in favor of heads or in favor of tails.”

From Griffin-Tversky (1992): Two possible *i.i.d.* coins, a  $\pi(h|\cdot) = \frac{3}{5}$  coin and a  $\pi(h|\cdot) = \frac{2}{5}$  coin, each with prob = .5. Observe a set of flips  $h, t$ . What is inferred?

Sample of (h,t)	Median $P(\theta = \frac{3}{5} h, t)$	Proper Bayes $B(\theta = \frac{3}{5} h, t)$
5,0	.92	.88
19,14	.60	.88
3,2	.60	.60
4,1	.80	.77
9,8	.54	.60
5,4	.55	.60
6,3	.67	.77
11,6	.64	.88
3,0	.85	.77
2,1	.63	.60
7,2	.77	.88
10,7	.60	.77

The Fundamental Theorem of Bayesian Inference about Which of Two Symmetrically Biased Coins have Generated A Sample of Heads and Tails: Proper inference in this setting is solely a function of  $h - t$ , and has nothing to do with  $h + t$ . Proof: Problem Set Maggie.

Proof: The Bayesian posterior probability that the bias ( $\Pr\{h\}$ ) is  $3/5$  given that you observe  $h$  heads and  $t$  tails is the probability that you observe  $h$  heads and  $t$  tails if the bias is  $3/5$  divided by the total probability that you observe  $h$  heads and  $t$  tails summed over both possible biases. Thus (canceling out the “number of ways we can get  $h$  heads” terms in the numerator and denominator),

posterior probability that  $\Pr\{h\} = 3/5 = [(3/5)^h (2/5)^t] / \{(3/5)^h (2/5)^t + (2/5)^h (3/5)^t\}$ .

Dividing the numerator and denominator through by  $[(3/5)^h (2/5)^t]$  gives

posterior probability that  $\Pr\{h\} = 3/5 = 1 / \{1 + (2/5)^{h-t} (3/5)^{t-h}\} = 1 / \{1 + (2/3)^{h-t}\}$ .

This formula yields the posteriors in the tables below.



Sample of (h,t)	$h - t$	Median $P(\theta = 3/5 h, t)$	Proper $B(\theta = \frac{3}{5} h, t)$
5,0	5	.92	.88
7,2	5	.77	.88
11,6	5	.64	.88
19,14	5	.60	.88
3,0	3	.85	.77
4,1	3	.80	.77
6,3	3	.67	.77
10,7	3	.60	.77
2,1	1	.63	.60
3,2	1	.60	.60
5,4	1	.55	.60
9,8	1	.54	.60

People instead seem to be basing their inference very largely on  $\frac{h}{h+t}$ .

Sample of (h,t)	% heads	Median $P(\theta = 3/5 h, t)$
5,0	100%	.92
3,0	100%	.85
4,1	80%	.80
7,2	78%	.77
6,3	67%	.67
2,1	67%	.63
11,6	65%	.60
3,2	60%	.60
10,7	59%	.60
19,14	58%	.60
5,4	55%	.55
9,8	53%	.54

Matthew's Claim: Only an egghead statistics geek would think that  $h-t$  tells you more about the coin's bias than  $\frac{h}{h+t}$ . Just because  $h-t$  is the correct statistic to use doesn't make it sensible or intuitive. Merely correct.

Sample of (h,t)	Median $P(\theta = 3/5)$	Bayes(3/5)
2,1	.63	.60
3,0	.85	.77
3,2	.60	.60
4,1	.80	.77
5,0	.92	.88
5,4	.55	.60
6,3	.67	.77
7,2	.77	.88
9,8	.54	.60
10,7	.60	.77
11,6	.64	.88
19,14	.60	.88

Relying so heavily on  $\frac{h}{h+t}$  in judging relative likelihoods, it might be (and in fact it is) that people think they are learning more from  $h - t$  in symmetric-coin inference problems than they are in small samples, and think they are learning less than they are in large samples.

More generally, *in inference problems*, people exhibit:

1. “Sample-Size Neglect”: Their inferences are too sensitive to  $\frac{h}{h+t}$  and too insensitive to  $h + t$ .
2. Non-Belief in the Law of Large Numbers—people don’t get just how much they learn from large data sets.
3. Belief in the Law of Small Numbers—people believe they are learning more from small samples than they really are.

## **Rabin on general modeling strategy:**

I believe it is very useful to emphasize the study of systematic departures from correct Bayesian information processing, always focusing on the Bayesian model as a base-line comparison. This is not the only reasonable approach to studying bounded rationality. But I think it makes a lot of sense because:

1. Bayesian rationality is what our models now assume. Note: A mini-literature in psychology responds to the K&T hypothesis that “In general, these heuristics are quite useful, but sometimes they lead to severe and systematic biases” with the rebuttal “No—while these heuristics may sometimes lead to severe and systematic biases, in general they are quite useful.” Irrespective of the merits of having such a debate in psychology, it is clear given the current status of economics that economists are more in need of understanding the 2<sup>nd</sup> half of the first sentence than the 2<sup>nd</sup> half of the second sentence. Economists haven’t omitted emphasis on the ways people are smart; we have assumed that people are inhumanly smart. The relevant insights for us are clearly identifying the severe and systematic biases.
2. This approach promotes emphasis on how probabilistic reasoning is not random or totally irrational. Central point to understand the heuristics-and-biases literature: Bounded Rationality = Randomness, and is a form of human rationality, not inhuman idiocy.
3. Bayesian updating is the unique normatively right way to process information.

4. Sticking as closely as possible to the Bayesian approach helps us take advantage of the existing apparatus in economics for thinking about statistical reasoning.

5. And it would be nice to follow the approach taken throughout the course and much of the psychology-and-economics literature: Model departures from “classical” models as contentions about parameter values in general models that embed the classical model as a special case. Embedding fully rational Bayesian updating as special cases of generalized models allows us to do good comparative statics and see where our results are coming from, and facilitates direct empirical testing.

### **Camerer on general modeling strategy**

The Bayesian approach is so simple and useful that it is hard to find equally simple formal alternatives consistent with Kahneman and Tversky’s heuristics. An appealing way to do so is to use the Bayesian framework but assume that people misspecify or misapply it in some way.

For example, Rabin and Schrag (*QJE* 1999) define “confirmatory bias” as the tendency to perceive data as more consistent with a prior hypothesis than they truly are; their model is otherwise fully Bayesian.

For example, Rabin (*QJE* 2002) models representativeness as the (mistaken) expectation that samples are drawn without replacement, and shows some fresh implications of that model (e.g, perceiving more skill among managers than truly exists).

## Modeling representativeness (Koszegi following Rabin (*QJE* 2002))

A person, “Freddy” in the paper, observes a sequence of binary signals of some quality, “good” or “bad.”

The signals are random, with a constant probability (“rate”) of being good, which Freddy has a correct Bayesian prior about.

Although the signals are really i.i.d., Freddy believes that they are generated by random draws without replacement from an “urn” of  $N$  signals, where the urn contains signals in proportions corresponding to the rate.

Other than misunderstanding the statistical process generating the outcomes, Freddy is completely rational: he always makes the correct inferences (and uses Bayes’ rule) given his wrong theory of the world.

$N \rightarrow \infty$  implies the person’s inferences approach Bayesian rationality;  $N \rightarrow 0$  implies they are more distorted.

The urn is completely replenished every 2 periods: A mathematical trick. (Convenient but less natural than continuous replenishment as in Rapoport and Budescu, “Randomization in Individual Choice Behavior,” *Psychological Review* (1997): People maintain window of  $k$  previous trials. Predict next trial  $t + 1$  will “balance” subsequence of  $k + 1$  (i.e., make relative proportions = probabilities).)

This model immediately (and trivially) yields the “gambler’s fallacy,” in that Freddy expects the second draw of a signal to be negatively correlated with the first.

The model also yields some insight into the hot hand fallacy and other puzzles. Virtually every sports fan believes that there is a systematic variation in the performance of each player from day to day, that the performance of a player during a particular period may be predictably better than expected on the basis of the player’s overall record. This was carefully tested, and rejected, for professional basketball by Gilovich, Vallone, and Tversky (1985 *Cognitive Psychology*). But people still believe it.

The hot hand fallacy at first seems the opposite of the gambler’s fallacy: The gambler’s fallacy is the belief that the next outcome is likely to be different from previous ones, whereas the hot hand fallacy is the belief that the outcome is likely to be similar to previous ones.

But both the gambler’s fallacy and the hot hand fallacy can be explained by this way of modeling representativeness. Intuitively, the gambler’s fallacy means people do not expect to see many streaks in a person’s performance. When they do see many streaks over a long period, the sequence will not feel representative of a completely independent random process. Hence, they will conclude that the person must have had a hot hand. This is the only way they can “rationally” explain the unexpected streaks they have observed.

First imagine that Freddy is playing roulette. He knows that the proportion of reds and blacks over time are equal.

Suppose  $N = 10$ , and that red came up twice in a row. That  $N = 10$  means that Freddy started off thinking that there were 5 red and 5 black balls in the run. With two red balls gone, Freddy expects the next draw to be black with probability  $5/8 > 1/2$ . That is, Freddy commits the gambler's fallacy. Intuitively, since he expects empirical proportions in small samples to resemble the true proportions, he expects initial imbalances to "correct themselves."

Now imagine that Freddy is trying to infer the skill level of a mutual-fund manager called Helga, where he does not know the probabilities of relevant outcomes. Freddy knows whether Helga has beaten the market average when running her managed fund zero times, once, or twice each of the last two quarters. He does not know Helga's skill level.

There is a probability  $q$  that Helga always beats the market, a probability  $q$  that she never beats the market, and a probability  $1 - 2q$  that she beats the market with probability  $1/2$ . Freddy has correct priors about the likelihood that Helga is skilled or unskilled.

In reality, a mediocre Helga's performance is independent from quarter to quarter.



Suppose that  $N = 2$ . What does Freddy think may happen? He correctly thinks that an unskilled Helga will fall short of the market both times and a skilled Helga will beat the market both times, but he thinks a mediocre Helga will beat the market exactly once and fall short of the market exactly once. As before, Freddy commits the gambler's fallacy: If he observes a mediocre Helga performing well, he thinks she is due for a bad performance.

What will Freddy infer after observing Helga having two good quarters in a row? He thinks only a skilled Helga can perform this well, so he will infer that Helga is skilled for sure. By the same logic, if Freddy he observes two bad performances in a row, he will conclude that Helga is unskilled for sure. Freddy overinfers Helga's skill from a small sample of extreme performances—he draws a more extreme conclusion than is justified by his observations.

This overinference is rooted in the same representativeness that is behind the gambler's fallacy. Since Freddy believes in the gambler's fallacy, he expects the mediocrity of Helga's decisions to be quickly reflected in a quarter of bad performance. Hence, he thinks that a couple of good performances indicate high skill.

Freddy's overinference regarding Helga's skill can lead him to several mistakes. First, he may be too eager to invest his money in Helga's fund if it has performed well recently. He will then be surprised when Helga's two good performances are followed (as is all too likely) by a bunch of average performances. That is, he underestimates "reversion to the mean".

If Helga follows her quarters of good performance with mediocre ones, Freddy may move his money pointlessly between investments.

Now consider what conclusion Freddy draws from a pair of mixed performances. He believes that only a mediocre Helga can have such performances, and he is in fact correct in this belief. Hence, Freddy overinfers only from extreme performances. Intuitively, the mistake Freddy makes is to underestimate the probability that a streak of good or bad performances arises purely by chance, so that he attributes extreme performances to skill. There is no similar mistake he makes with mixed performances.

Freddy's overinference about Helga's skill from a short sequence of good performances is exactly like a false belief in the hot hand: He too easily concludes that Helga's skill is "hot." But this can only happen because Freddy thinks it is a priori possible that Helga is a skilled manager.

However, the model implies that Freddy will come to believe this even if it is not the case: Upon observing the world, he will come to believe that there is variation in manager skill when in fact there is none.

Suppose Freddy does not have a good idea about the distribution of skills among mutual-fund managers, and is trying to learn this from observing the performance of mutual funds.

Consider an extreme case: in reality all mutual-fund managers are mediocre, beating the market each quarter with probability one-half.

Freddy observes two quarters of performance by a large number of mutual funds. Since in reality all managers are mediocre and have independent performances, Freddy will observe that one-quarter of the mutual funds beat the market twice, half beat the market once, and one-quarter do worse than the market twice. But Freddy falsely believes that only skilled mutual-fund managers can beat the market twice in a row, so he concludes that one-quarter of the mutual-fund managers are skilled. Similarly, he concludes that one-quarter of the mutual-fund managers are unskilled. That is, he overestimates the variation in skill in the population.

Intuitively, since Freddy believes in the gambler's fallacy, he thinks streaks of good or bad performance are unlikely. So if he watches several analysts for two quarters, because he underestimates how often average analysts will have consecutive successful or unsuccessful years, he interprets what he sees as evidence of the existence of good and bad analysts.

This sort of analysis can be done with tedious algebra for less extreme cases (Rabin):

Let good performances be  $a$ , and bad  $b$ . Then:

	True Probability			Freddy's Predicted Probability		
	$aa$	$ab$ (or $ba$ )	$bb$	$aa$	$ab$ (or $ba$ )	$bb$
If $\theta = \frac{3}{4}$	$\frac{9}{16}$	$\frac{6}{16}$	$\frac{1}{16}$	$\frac{3}{6}$	$\frac{3}{6}$	0
If $\theta = \frac{1}{2}$	$\frac{4}{16}$	$\frac{8}{16}$	$\frac{4}{16}$	$\frac{1}{6}$	$\frac{4}{6}$	$\frac{1}{6}$
If $\theta = \frac{1}{4}$	$\frac{1}{16}$	$\frac{6}{16}$	$\frac{9}{16}$	0	$\frac{3}{6}$	$\frac{3}{6}$

Claim: Freddy will over-infer likelihood of  $\theta = \frac{3}{4}$  upon observing  $aa$ . Exercise: Calculate his and a Bayesian's posterior beliefs about Helga following  $aa$ .

Given his priors  $q$ , with what chance does Freddy expect to see  $aa$ ,  $ab$ , and  $bb$  performances by Helga?

Freddy believes the likelihood of  $aa$  is  $q\frac{3}{6} + (1 - 2q)\frac{1}{6} + q0 = \frac{1}{6} + \frac{1}{6}q$   
 Freddy believes the likelihood of  $ab$  is  $q\frac{3}{6} + (1 - 2q)\frac{4}{6} + q\frac{3}{6} = \frac{32}{48} - \frac{16}{48}q$   
 Freddy believes the likelihood of  $bb$  is  $q0 + (1 - 2q)\frac{1}{6} + q\frac{3}{6} = \frac{1}{6} + \frac{1}{6}q$ .

Actual likelihood of  $aa$  is  $q\frac{9}{16} + (1 - 2q)\frac{4}{16} + q\frac{1}{16} = \frac{4}{16} + \frac{2}{16}q$   
 Actual likelihood of  $ab$  is  $q\frac{6}{16} + (1 - 2q)\frac{8}{16} + q\frac{6}{16} = \frac{24}{48} - \frac{12}{48}q$   
 Actual likelihood of  $bb$  is  $q\frac{1}{16} + (1 - 2q)\frac{4}{16} + q\frac{9}{16} = \frac{4}{16} + \frac{2}{16}q$

From Rabin (*QJE* 2002)

#### IV. OVERINFERENCE

The most interesting implications of the law of small numbers come when Freddy is uncertain about the true rate, and makes inferences about the rate from the signals he observes. Suppose, for instance, that an observer believes that there is an equal chance a fund manager can be any of three types, bad, average, or good, who outperforms other mutual funds  $1/4$ ,  $1/2$ , or  $3/4$  of the time, respectively. What does he infer from two successful years in a row by a particular fund? A Bayesian thinks such a sequence occurs with probability  $1/4 \cdot 1/4 = 1/16$  for bad funds,  $2/4 \cdot 2/4 = 4/16$  for average funds, and  $3/4 \cdot 3/4 = 9/16$  for good funds. But an  $N = 4$ -Freddy believes that the probabilities are  $1/4 \cdot 0/3 = 0/12$  if the fund is bad,  $2/4 \cdot 1/3 = 2/12$  if the fund is average, and  $3/4 \cdot 2/3 = 6/12$  if the fund is good. For each rate, Freddy assigns a lower probability to a streak of two  $a$ 's than a Bayesian assigns—because he believes that no matter the rate, drawing the first  $a$  means there are fewer  $a$ 's left for the second draw. But more importantly, Freddy's beliefs are too skewed toward believing that the fund is good, since making one less  $a$  available for the second draw has a proportionately greater impact when there are fewer  $a$ 's to begin with. From his priors, Freddy forms probabilistic beliefs about the rate given an observed sequence of signals using a sort of warped Bayes' Law—applying Bayes' Law with his mistaken beliefs about how likely each sequence is given an underlying rate. While a Bayesian believes the probability that the analyst is good is  $18/28$ , Freddy believes that the probability is  $21/28 > 18/28$ .

By contrast, if Freddy observed a large number of signals from each of several different sources, then he is too likely to believe that the rate is less extreme than it is.

This is because he struggles to explain why he is observing so many streaks of rare signals, which he thinks are very unlikely.

To explain such streaks, he may come to believe that the true rate is close to 50/50, even if this does not accord with the overall frequency of the signals, assuming that there is underlying variation even when there is none.

Such a false belief corresponds to the hot hand fallacy.

See Rabin and Vayanos, “The Gambler's and Hot-Hand Fallacies in a Dynamic-Inference Model”.

## **Further Issues (not covered in class)**

### **Overconfidence**

98% confidence interval only captures 60% of the distribution.  
100% is actually 80% and 0% is actually 20%

### **Optimism/Wishful Thinking**

Unrealistic view of personal abilities/prospects  
90% of drivers claim above average skill  
99% of freshman claim superior intelligence

### **Confirmatory Bias**

People selectively either ignore or misread further (ambiguous) information as supporting initial hypotheses. That is, once a person has formed a strong hypothesis about what is true, she tends to selectively ignore contrary evidence and carefully consider supporting evidence.

### **Anchoring**

People often try to answer a question by starting at some first-pass guess based on memory or the environment, and then adjusting that guess until they are satisfied with the answer. Even after the adjustment, people's judgment seems to be colored by their original guess or anchor.

## **Availability Biases**

People often assess the frequency of a class or the probability of an event by the ease with which instances or occurrences can be brought to mind. To answer questions such as “What percentage of commercial flights crash per year?”; “What is the risk of heart attack among middle-aged men?”; or “Are there more suicides or more homicides in the United States each year?”, most people (who do not know the answer) try to recall instances of plane crashes, heart attacks, suicides, or murders they have heard about from acquaintances or in the news. The easier they can recall instances of the event, the more likely they perceive it to be.

## **Curse of Knowledge**

The hindsight bias is closely related to the curse of knowledge. People cannot abandon their own perspective, even when they know others are in a different situation, and they are highly motivated to communicate well.

## **Incomplete Debiasing**

Suppose a person is told that A is true, where A leads to some conclusion X. Then, she is told that A is actually not true—it was a mistake—and she believes this. Despite this, she will believe in X more than if she never heard A.