

writemac

THINKING, FAST AND SLOW

BY

DANIEL KAHNEMAN

Notes

The title of the book refers to two modes of thinking, which he refers to as:

* “System 1” = The instant, unconscious, automatic, emotional, intuitive thinking.

* “System 2” = The slower, conscious, rational, reasoning, deliberate thinking.

=====

EXPERTISE:

Expert intuition: The situation has provided a cue; this cue has given the expert access to information stored in memory, and the information provides the answer. Intuition is nothing more and nothing less than recognition.

Valid intuitions develop when experts have learned to recognize familiar elements in a new situation and to act in a manner that is appropriate to it.

Philip Tetlock's book "Expert Political Judgment: How Good Is It? How Can We Know?" - gathered more than 80,000 predictions. The experts performed worse than they would have if they had simply assigned equal probabilities. Even in the region they knew best, experts were not significantly better than nonspecialists.

People who spend their time, and earn their living, studying a particular topic produce poorer predictions than dart-throwing monkeys.

Those with the most knowledge are often less reliable. The reason is that the person who acquires more knowledge develops an enhanced illusion of her skill and becomes unrealistically overconfident.

Hedgehogs "know one big thing" and have a theory about the world; they account for particular events within a coherent framework, bristle with impatience toward those who don't see things their way, and are confident in their forecasts. They are also especially reluctant to admit error.

It is much easier to strive for perfection when you are never bored.

Flow neatly separates the two forms of effort: concentration on the task and the deliberate control of attention.

In a state of flow, maintaining focused attention on these absorbing activities requires no exertion of self-control, thereby freeing resources to be directed to the task at hand.

Many people are overconfident, prone to place too much faith in their intuitions. They apparently find cognitive effort at least mildly unpleasant and avoid it as much as possible.

Putting the participants in a good mood before the test by having them think happy thoughts more than doubled accuracy. An even more striking result is that unhappy subjects were completely

incapable of performing the intuitive task accurately; their guesses were no better than random. Mood evidently affects the operation of System 1: when we are uncomfortable and unhappy, we lose touch with our intuition.

When in a good mood, people become more intuitive and more creative but also less vigilant and more prone to logical errors. Here again, as in the mere exposure effect, the connection makes biological sense. A good mood is a signal that things are generally going well, the environment is safe, and it is all right to let one's guard down. A bad mood indicates that things are not going very well, there may be a threat, and vigilance is required.

Surprise itself is the most sensitive indication of how we understand our world and what we expect from it.

The main function of System 1 is to maintain and update a model of your personal world, which represents what is normal in it.

When System 2 is otherwise engaged, we will believe almost anything. System 1 is gullible and biased to believe, System 2 is in charge of doubting and unbelieving, but System 2 is sometimes busy, and often lazy.

Understanding a statement must begin with an attempt to believe it: you must first know what the idea would mean if it were true. Only then can you decide whether or not to unbeliev it. The initial attempt to believe is an automatic operation of System 1.

Unbelieving is an operation of System 2.

The operations of associative memory contribute to a general confirmation bias. When asked, "Is Sam friendly?" different instances of Sam's behavior will come to mind than would if you had been asked "Is Sam unfriendly?" A deliberate search for confirming evidence, known as positive test strategy, is also how System 2 tests a hypothesis. Contrary to the rules of philosophers of science, who

advise testing hypotheses by trying to refute them, people (and scientists, quite often) seek data that are likely to be compatible with the beliefs they currently hold. The confirmatory bias of System 1 favors uncritical acceptance of suggestions and exaggeration of the likelihood of extreme and improbable events.

Herbert Simon's definition of intuition: Expertise in a domain is not a single skill but rather a large collection of miniskills.

The confidence that people have in their intuitions is not a reliable guide to their validity. In other words, do not trust anyone - including yourself - to tell you how much you should trust their judgment.

When do judgments reflect true expertise?

An environment that is sufficiently regular to be predictable an opportunity to learn these regularities through prolonged practice. When both these conditions are satisfied, intuitions are likely to be skilled.

Intuition cannot be trusted in the absence of stable regularities in the environment.

If the environment is sufficiently regular and if the judge has had a chance to learn its regularities, the associative machinery will recognize situations and generate quick and accurate predictions and decisions. You can trust someone's intuitions if these conditions are met.

When evaluating expert intuition you should always consider whether there was an adequate opportunity to learn the cues, even in a regular environment.

"Does he really believe that the environment of start-ups is sufficiently regular to justify an intuition that goes against the base rates?"

“Did he really have an opportunity to learn? How quick and how clear was the feedback he received on his judgments?”

The proper way to elicit information from a group is not by starting with a public discussion but by confidentially collecting each person’s judgment.

=====

JUMPING TO CONCLUSIONS

The way to block errors that originate in System 1 is simple in principle: recognize the signs that you are in a cognitive minefield, slow down, and ask for reinforcement from System 2.

When you see lines with fins pointing in different directions, you will recognize the situation as one in which you should not trust your impressions of length. Unfortunately, this sensible procedure is least likely to be applied when it is needed most.

Organizations are better than individuals when it comes to avoiding errors, because they naturally think more slowly and have the power to impose orderly procedures. Organizations can institute and enforce the application of useful checklists,

When the question is difficult and a skilled solution is not available, intuition still has a shot: an answer may come to mind quickly - but it is not an answer to the original question.

When faced with a difficult question, we often answer an easier one instead, usually without noticing the substitution.

An easy question (How do I feel about it?) serves as an answer to a much harder question (What do I think about it?).

System 1 effortlessly originates impressions and feelings that are the main sources of the explicit beliefs and deliberate choices of System 2.

The automatic operations of System 1 generate surprisingly complex

patterns of ideas, but only the slower System 2 can construct thoughts in an orderly series of steps.

You can also feel a surge of conscious attention whenever you are surprised. System 2 is activated when an event is detected that violates the model of the world that System 1 maintains.

Most of what you (your System 2) think and do originates in your System 1, but System 2 takes over when things get difficult, and it normally has the last word.

Continuous vigilance is not necessarily good, and it is certainly impractical. Constantly questioning our own thinking would be impossibly tedious, and System 2 is much too slow and inefficient to serve as a substitute for System 1 in making routine decisions. The best we can do is a compromise: learn to recognize situations in which mistakes are likely and try harder to avoid significant mistakes

Anything that occupies your working memory reduces your ability to think.

Test Questions were chosen because they also invite an intuitive answer that is both compelling and wrong:

Students who scored very low on this test - their supervisory function of System 2 is weak - and they are prone to answer questions with the first idea that comes to mind and unwilling to invest the effort needed to check their intuitions.

Individuals who uncritically follow their intuitions about puzzles are also prone to accept other suggestions from System 1. In particular, they are impulsive, impatient, and keen to receive immediate gratification.

What makes some people more susceptible than others to biases of judgment? Stanovich published his conclusions in a book titled *Rationality and the Reflective Mind*.

Superficial or “lazy” thinking is a flaw in the reflective mind, a failure of rationality.

Rationality should be distinguished from intelligence.

When information is scarce, which is a common occurrence, System 1 operates as a machine for jumping to conclusions. You did not start by asking, “What would I need to know before I formed an opinion about the quality of someone’s leadership?” System 1 got to work on its own from the first adjective.

The combination of a coherence-seeking System 1 with a lazy System 2 implies that System 2 will endorse many intuitive beliefs, which closely reflect the impressions generated by System 1.

Based on brief exposure to photographs and without any political context: In about 70% of the races for senator, congressman, and governor, the election winner was the candidate whose face had earned a higher rating of competence.

A remarkable aspect of your mental life is that you are rarely stumped. You have intuitive feelings and opinions about almost everything that comes your way. You often have answers to questions that you do not completely understand, relying on evidence that you can neither explain nor defend.

If a satisfactory answer to a hard question is not found quickly, System 1 will find a related question that is easier and will answer it.

If they had been given indefinite time and told to follow logic and not to answer until they were sure of their answer, I believe that most of our subjects would have avoided the conjunction fallacy. However, their vacation did not depend on a correct answer; they spent very little time on it, and were content to answer as if they had only been “asked for their opinion.” The laziness of System 2 is an important fact of life.

Following our intuitions is more natural, and somehow more pleasant, than acting against them.

You cannot help dealing with the limited information you have as if it were all there is to know. You build the best possible story from the information available to you, and if it is a good story, you believe it. Paradoxically, it is easier to construct a coherent story when you know little, when there are fewer pieces to fit into the puzzle. Our comforting conviction that the world makes sense rests on a secure foundation: our almost unlimited ability to ignore our ignorance.

Poor evidence can make a very good story.

For some of our most important beliefs we have no evidence at all, except that people we love and trust hold these beliefs.

Cognitive illusions can be more stubborn than visual illusions. What you learned about the Müller-Lyer illusion did not change the way you see the lines.

Intuition adds value, but only after a disciplined collection of objective information and disciplined scoring of separate traits.

Do not simply trust intuitive judgment - your own or that of others - but do not dismiss it, either.

=====

STATISTICS

People are prone to apply causal thinking inappropriately, to situations that require statistical reasoning. Statistical thinking derives conclusions about individual cases from properties of categories and ensembles. Unfortunately, System 1 does not have the capability for this mode of reasoning; System 2 can learn to think statistically, but few people receive the necessary training.

From the same urn, two very patient marble counters take turns. Jack draws 4 marbles on each trial, Jill draws 7. They both record

each time they observe a homogeneous sample - all white or all red. If they go on long enough, Jack will observe such extreme outcomes more often than Jill - by a factor of 8 (the expected percentages are 12.5% and 1.56%). Again, no hammer, no causation, but a mathematical fact: samples of 4 marbles yield extreme results more often than samples of 7 marbles do. Now imagine the population of the United States as marbles in a giant urn. Some marbles are marked KC, for kidney cancer. You draw samples of marbles and populate each county in turn. Rural samples are smaller than other samples. Just as in the game of Jack and Jill, extreme outcomes (very high and/or very low cancer rates) are most likely to be found in sparsely populated counties. This is all there is to the story.

People should regard their statistical intuitions with proper suspicion and replace impression formation by computation whenever possible.

We are prone to exaggerate the consistency and coherence of what we see. The exaggerated faith of researchers in what can be learned from a few observations is closely related to the halo effect, the sense we often get that we know and understand a person about whom we actually know very little.

The associative machinery seeks causes. The difficulty we have with statistical regularities is that they call for a different approach. Instead of focusing on how the event at hand came to be, the statistical view relates it to what could have happened instead. Nothing in particular caused it to be what it is - chance selected it from among its alternatives. Our predilection for causal thinking exposes us to serious mistakes in evaluating the randomness of truly random events.

Bad schools also tend to be smaller than average. The truth is that small schools are not better on average; they are simply more variable.

A basic limitation in the ability of our mind to deal with small risks: we either ignore them altogether or give them far too much weight - nothing in between.

We tend to overweight small risks and are willing to pay far more than expected value to eliminate them altogether.

When an unlikely event becomes the focus of attention, we will assign it much more weight than its probability deserves.

Reducing or mitigating the risk is not adequate; to eliminate the worry the probability must be brought down to zero.

People overestimate the probabilities of unlikely events. People overweight unlikely events in their decisions.

What is the probability that a baby born in your local hospital will be released within three days? You were asked to estimate the probability of the baby going home, but you almost certainly focused on the events that might cause a baby not to be released within the normal period. Our mind has a useful capability to focus spontaneously on whatever is odd, different, or unusual.

The unlikely event became focal.

Your estimate of the frequency of problems was too high.

The successful execution of a plan is specific and easy to imagine when one tries to forecast the outcome of a project. In contrast, the alternative of failure is diffuse, because there are innumerable ways for things to go wrong. Entrepreneurs and the investors who evaluate their prospects are prone both to overestimate their chances and to overweight their estimates.

The decision weight for a 90% chance was 71.2 and the decision weight for a 10% chance was 18.6.

The valuation of gambles was much less sensitive to probability when the (fictitious) outcomes were emotional, than when the outcomes were gains or losses of cash.

The fear of an impending electric shock was essentially uncorrelated with the probability of receiving the shock. The mere possibility of a shock triggered the full-blown fear response.

Urn A contains 10 marbles, of which 1 is red.

Urn B contains 100 marbles, of which 8 are red.

30%–40% of students choose the urn with the larger number of winning marbles, rather than the urn that provides a better chance of winning.

If your attention is drawn to the winning marbles, you do not assess the number of nonwinning marbles with the same care. Vivid imagery contributes to denominator neglect.

Low-probability events are much more heavily weighted when described in terms of relative frequencies (how many) than when stated in more abstract terms of “chances,” “risk,” or “probability” (how likely). As we have seen, System 1 is much better at dealing with individuals than categories.

“A disease that kills 1,286 people out of every 10,000” was judged more dangerous than a disease that “kills 24.4 out of 100.”

=====

PRIMING / ANCHORING

Your actions and your emotions can be primed by events of which you are not even aware.

The common admonition to “act calm and kind regardless of how you feel” is very good advice: You are likely to be rewarded by actually feeling calm and kind.

Money-primed people become more independent than they would be without the associative trigger. They persevered almost twice as long in trying to solve a very difficult problem before they asked the experimenter for help, a crisp demonstration of increased self-reliance. Money-primed people are also more selfish: they were much less willing to spend time helping another student who pretended to be confused about an experimental task. When an experimenter clumsily dropped a bunch of pencils on the floor, the participants with money (unconsciously) on their mind picked up fewer pencils.

Money-primed undergraduates also showed a greater preference for being alone. The general theme of these findings is that the idea of money primes individualism: a reluctance to be involved with others, to depend on others, or to accept demands from others. The psychologist who has done this remarkable research, Kathleen Vohs,

Living in a culture that surrounds us with reminders of money may shape our behavior and our attitudes in ways that we do not know about and of which we may not be proud. Some cultures provide frequent reminders of respect, others constantly remind their members of God, and some societies prime obedience by large images of the Dear Leader.

Reminding people of their mortality increases the appeal of authoritarian ideas, which may become reassuring in the context of the terror of death.

Feeling that one's soul is stained appears to trigger a desire to cleanse one's body.

A reliable way to make people believe in falsehoods is frequent repetition, because familiarity is not easily distinguished from truth.

Words that were presented more frequently were rated much more favorably than the words that had been shown only once or twice.

Biological fact: an organism should react cautiously to a novel stimulus, with withdrawal and fear.

If repeated exposure of a stimulus is followed by nothing bad, such a stimulus will eventually become a safety signal.

The evaluation of the risk depends on the choice of a measure - with the obvious possibility that the choice may have been guided by a preference for one outcome or another. He goes on to conclude that "defining risk is thus an exercise in power."

=====

AVAILABILITY HEURISTIC

The experience of familiarity has a simple but powerful quality of 'pastness' that seems to indicate that it is a direct reflection of prior experience.

This quality of pastness is an illusion.

A name you've seen before will look familiar when you see it because you will see it more clearly. Words that you have seen before become easier to see again - you can identify them better than other words when they are shown very briefly or masked by noise, and you will be quicker (by a few hundredths of a second) to read them than to read other words. In short, you experience greater cognitive ease in perceiving a word you have seen earlier, and it is this sense of ease that gives you the impression of familiarity.

Information that is not retrieved (even unconsciously) from memory might as well not exist. System 1 excels at constructing the best possible story that incorporates ideas currently activated, but it does not (cannot) allow for information it does not have. The measure of success for System 1 is the coherence of the story it manages to create.

Availability heuristic, like other heuristics of judgment, substitutes one question for another: you wish to estimate the size of a category

or the frequency of an event, but you report an impression of the ease with which instances come to mind.

Discover how the heuristic leads to biases by following a simple procedure: list factors other than frequency that make it easy to come up with instances. Each factor in your list will be a potential source of bias.

People are less confident in a choice when they are asked to produce more arguments to support it.

Students who listed more ways to improve the class rated it higher!

Death by accidents was judged to be more than 300 times more likely than death by diabetes, but the true ratio is 1:4. The lesson is clear: estimates of causes of death are warped by media coverage. The coverage is itself biased toward novelty and poignancy.

The ease with which ideas of various risks come to mind and the emotional reactions to these risks are inextricably linked. Frightening thoughts and images occur to us with particular ease, and thoughts of danger that are fluent and vivid exacerbate fear.

Anchoring effect. It occurs when people consider a particular value for an unknown quantity before estimating that quantity. What happens is one of the most reliable and robust results of experimental psychology: the estimates stay close to the number that people considered.

The same house will appear more valuable if its listing price is high than if it is low, even if you are determined to resist the influence of this number.

Any number that you are asked to consider as a possible solution to an estimation problem will induce an anchoring effect.

On some days, a sign on the shelf said limit of 12 per person. On other days, the sign said no limit per person. Shoppers purchased an

average of 7 cans when the limit was in force, twice as many as they bought when the limit was removed. Anchoring is not the sole explanation. Rationing also implies that the goods are flying off the shelves, and shoppers should feel some urgency about stocking up.

=====

STEREOTYPES

Stereotypes are statements about the group that are (at least tentatively) accepted as facts about every member. Here are two examples: Most of the graduates of this inner-city school go to college. Interest in cycling is widespread in France.

You will be reminded of these facts when you think about the likelihood that a particular graduate of the school will attend college, or when you wonder whether to bring up the Tour de France in a conversation with a Frenchman you just met. Stereotyping is a bad word in our culture, but in my usage it is neutral. One of the basic characteristics of System 1 is that it represents categories as norms and prototypical exemplars.

In sensitive social contexts, we do not want to draw possibly erroneous conclusions about the individual from the statistics of the group. We consider it morally desirable for base rates to be treated as statistical facts about the group rather than as presumptive facts about individuals. In other words, we reject causal base rates. The social norm against stereotyping, including the opposition to profiling, has been highly beneficial in creating a more civilized and more equal society. It is useful to remember, however, that neglecting valid stereotypes inevitably results in suboptimal judgments. Resistance to stereotyping is a laudable moral position, but the simplistic idea that the resistance is costless is wrong. The costs are worth paying to achieve a better society, but denying that the costs exist, while satisfying to the soul and politically correct, is not scientifically defensible.

=====

REVERT TO THE MEAN

When you have doubts about the quality of the evidence: let your judgments of probability stay close to the base rate.

How to Discipline Intuition:

You should not let yourself believe whatever comes to your mind. To be useful, your beliefs should be constrained by the logic of probability.

Base rates matter, even in the presence of evidence about the case at hand.

Intuitive impressions of the diagnosticity of evidence are often exaggerated.

The combination of WYSIATI and associative coherence tends to make us believe in the stories we spin for ourselves.

Anchor your judgment of the probability of an outcome on a plausible base rate. Question the diagnosticity of your evidence.

“This start-up looks as if it could not fail, but the base rate of success in the industry is extremely low. How do we know this case is different?”

“They keep making the same mistake: predicting rare events from weak evidence. When the evidence is weak, one should stick with the base rates.”

An important principle of skill training: rewards for improved performance work better than punishment of mistakes.

A significant fact of the human condition: the feedback to which life exposes us is perverse. Because we tend to be nice to other people

when they please us and nasty when they do not, we are statistically punished for being nice and rewarded for being nasty.

The “Sports Illustrated jinx,” the claim that an athlete whose picture appears on the cover of the magazine is doomed to perform poorly the following season. Overconfidence and the pressure of meeting high expectations are often offered as explanations. But there is a simpler account of the jinx: an athlete who gets to be on the cover of Sports Illustrated must have performed exceptionally well in the preceding season, probably with the assistance of a nudge from luck - and luck is fickle.

If you treated a group of depressed children for some time with an energy drink, they would show a clinically significant improvement. It is also the case that depressed children who spend some time standing on their head or hug a cat for twenty minutes a day will also show improvement. Most readers of such headlines will automatically infer that the energy drink or the cat hugging caused an improvement, but this conclusion is completely unjustified. Depressed children are an extreme group, they are more depressed than most other children - and extreme groups regress to the mean over time. The correlation between depression scores on successive occasions of testing is less than perfect, so there will be regression to the mean: depressed children will get somewhat better over time even if they hug no cats and drink no Red Bull. In order to conclude that an energy drink - or any other treatment - is effective, you must compare a group of patients who receive this treatment to a “control group” that receives no treatment (or, better, receives a placebo). The control group is expected to improve by regression alone, and the aim of the experiment is to determine whether the treated patients improve more than regression can explain.

“She says experience has taught her that criticism is more effective than praise. What she doesn’t understand is that it’s all due to regression to the mean.”

“Perhaps his second interview was less impressive than the first

because he was afraid of disappointing us, but more likely it was his first that was unusually good.”

The basic message of Built to Last and other similar books is that good managerial practices can be identified and that good practices will be rewarded by good results. Both messages are overstated. The comparison of firms that have been more or less successful is to a significant extent a comparison between firms that have been more or less lucky. Knowing the importance of luck, you should be particularly suspicious when highly consistent patterns emerge from the comparison of successful and less successful firms. In the presence of randomness, regular patterns can only be mirages.

On average, the gap in corporate profitability and stock returns between the outstanding firms and the less successful firms studied in Built to Last shrank to almost nothing in the period following the study.

A study of Fortune’s “Most Admired Companies” finds that over a twenty-year period, the firms with the worst ratings went on to earn much higher stock returns than the most admired firms.

The average gap must shrink, because the original gap was due in good part to luck, which contributed both to the success of the top firms and to the lagging performance of the rest. We have already encountered this statistical fact of life: regression to the mean.

=====

PREDICTIONS

Philip Tetlock's book "Expert Political Judgment: How Good Is It? How Can We Know?" - gathered more than 80,000 predictions. The experts performed worse than they would have if they had simply assigned equal probabilities. Even in the region they knew best, experts were not significantly better than nonspecialists.

People who spend their time, and earn their living, studying a particular topic produce poorer predictions than dart-throwing monkeys.

“Miswanting”: bad choices that arise from errors of affective forecasting.

Eliminating redundancy from your sources of information is always a good idea.

The magic of error reduction works well only when the observations are independent and their errors uncorrelated. If the observers share a bias, the aggregation of judgments will not reduce it. Allowing the observers to influence each other effectively reduces the size of the sample, and with it the precision of the group estimate. To derive the most useful information from multiple sources of evidence, you should always try to make these sources independent of each other.

The prediction of the future is not distinguished from an evaluation of current evidence - prediction matches evaluation.

People are asked for a prediction but they substitute an evaluation of the evidence, without noticing that the question they answer is not the one they were asked. This process is guaranteed to generate predictions that are systematically biased; they completely ignore regression to the mean.

Start with an estimate of average GPA. Determine the GPA that matches your impression of the evidence. Estimate the correlation between your evidence and GPA. If the correlation is .30, move 30% of the distance from the average to the matching GPA.

Suppose that I predict for each golfer in a tournament that his score on day 2 will be the same as his score on day 1. This prediction does not allow for regression to the mean: the golfers who fared well on day 1 will on average do less well on day 2, and those who did poorly will mostly improve. When they are eventually compared to actual

outcomes, nonregressive predictions will be found to be biased. They are on average overly optimistic for those who did best on the first day and overly pessimistic for those who had a bad start.

Similarly, if you use childhood achievements to predict grades in college without regressing your predictions toward the mean, you will more often than not be disappointed by the academic outcomes of early readers and happily surprised by the grades of those who learned to read relatively late. The corrected intuitive predictions eliminate these biases.

A baseline prediction, which you would make if you knew nothing about the case at hand. In the categorical case, it was the base rate. In the numerical case, it is the average outcome in the relevant category.

An intuitive prediction, which expresses the number that comes to your mind.

Aim for a prediction that is intermediate between the baseline and your intuitive response.

In the default case of no useful evidence, you stay with the baseline.

Find some reason to doubt that the correlation between your intuitive judgment and the truth is perfect, and you will end up somewhere between the two poles.

Intuitive predictions tend to be overconfident and overly extreme.

Correcting your intuitions may complicate your life.

Unbiased predictions permit the prediction of rare or extreme events only when the information is very good. If you expect your predictions to be of modest validity, you will never guess an outcome that is either rare or far from the mean. If your predictions are unbiased, you will never have the satisfying experience of correctly calling an extreme case. You will never be able to say, "I thought so!"

The ultimate test of an explanation is whether it would have made the event predictable in advance.

We can know something only if it is both true and knowable. But the crisis was not knowable. What is perverse about the use of know in this context is not that some individuals get credit for prescience that they do not deserve. It is that the language implies that the world is more knowable than it is. It helps perpetuate a pernicious illusion.

=====

CAUSAL

Our mind is strongly biased toward causal explanations.

Students “quietly exempt themselves” (and their friends and acquaintances) from the conclusions of experiments that surprise them.

When they presented their students with a surprising statistical fact, the students managed to learn nothing at all. But when the students were surprised by individual cases - two nice people who had not helped - they immediately made the generalization.

Subjects’ unwillingness to deduce the particular from the general was matched only by their willingness to infer the general from the particular. This is a profoundly important conclusion. People who are taught surprising statistical facts about human behavior may be impressed to the point of telling their friends about what they have heard, but this does not mean that their understanding of the world has really changed. The test of learning psychology is whether your understanding of situations you encounter has changed, not whether you have learned a new fact.

Surprising individual cases have a powerful impact and are a more effective tool for teaching psychology because the incongruity must be resolved and embedded in a causal story.

You are more likely to learn something by finding surprises in your own behavior than by hearing surprising facts about people in general.

When our attention is called to an event, associative memory will look for its cause. Any cause that is already stored in memory. Causal explanations will be evoked when regression is detected, but they will be wrong because the truth is that regression to the mean has an explanation but does not have a cause.

The explanatory stories that people find compelling are simple; are concrete rather than abstract; assign a larger role to talent, stupidity, and intentions than to luck; and focus on a few striking events that happened rather than on the countless events that failed to happen.

=====
RANDOMNESS

Success = talent + luck
great success = a little more talent + a lot of luck.

The idea that the future is unpredictable is undermined every day by the ease with which the past is explained.

The idea that large historical events are determined by luck is profoundly shocking, although it is demonstrably true.

The line that separates the possibly predictable future from the unpredictable distant future is yet to be drawn.

=====
INVESTING

Absence of bias is not always what matters most.

A venture capitalist will never be told that the probability of success for a start-up in its early stages is “very high.”

When a venture capitalist looks for “the next big thing,” the risk of missing the next Google or Facebook is far more important than the risk of making a modest investment in a start-up that ultimately fails. The goal of venture capitalists is to call the extreme cases correctly, even at the cost of overestimating the prospects of many other ventures.

Some of us may need the security of distorted estimates to avoid paralysis. If you choose to delude yourself by accepting extreme predictions, however, you will do well to remain aware of your self-indulgence. Perhaps the most valuable contribution of the corrective procedures I propose is that they will require you to think about how much you know.

I have heard of too many people who “knew well before it happened that the 2008 financial crisis was inevitable.”

When an unpredicted event occurs, we immediately adjust our view of the world to accommodate the surprise.

We have an imperfect ability to reconstruct past states of knowledge.

Once you adopt a new view of the world (or of any part of it), you immediately lose much of your ability to recall what you used to believe before your mind changed.

The illusion that one has understood the past feeds the further illusion that one can predict and control the future. These illusions are comforting. They reduce the anxiety that we would experience if we allowed ourselves to fully acknowledge the uncertainties of existence. We all have a need for the reassuring message that actions have appropriate consequences, and that success will reward wisdom and courage. Many business books are tailor-made to satisfy this need.

“When you sell a stock,” I asked, “who buys it?” He answered with a wave in the vague direction of the window, indicating that he expected the buyer to be someone else very much like him. That was odd: What made one person buy and the other sell? What did the sellers think they knew that the buyers did not?

A major industry appears to be built largely on an illusion of skill.

The buyers and sellers know that they have the same information; they exchange the stocks primarily because they have different opinions.

If all assets in a market are correctly priced, no one can expect either to gain or to lose by trading. Perfect prices leave no scope for cleverness, but they also protect fools from their own folly.

For the large majority of individual investors, taking a shower and doing nothing would have been a better policy than implementing the ideas that came to their minds.

Individual investors predictably flock to companies that draw their attention because they are in the news. Professional investors are more selective in responding to news. These findings provide some justification for the label of “smart money” that finance professionals apply to themselves.

A basic test of skill: persistent achievement. The diagnostic for the existence of any skill is the consistency of individual differences in achievement.

The illusion of skill is not only an individual aberration; it is deeply ingrained in the culture of the industry. Facts that challenge such basic assumptions - and thereby threaten people’s livelihood and self-esteem - are simply not absorbed. The mind does not digest them.

Skill in evaluating the business prospects of a firm is not sufficient for successful stock trading, where the key question is whether the information about the firm is already incorporated in the price of its stock. Traders apparently lack the skill to answer this crucial question, but they appear to be ignorant of their ignorance.

Large numbers of individuals in that world believe themselves to be among the chosen few who can do what they believe others cannot.

The financial benefits of self-employment are mediocre: given the same qualifications, people achieve higher average returns by selling their skills to employers than by setting out on their own. The evidence suggests that optimism is widespread, stubborn, and costly.

Daniel Bernoulli argued that a gift of 10 ducats has the same utility to someone who already has 100 ducats as a gift of 20 ducats to someone whose current wealth is 200 ducats.

The psychological response to a change of wealth is inversely proportional to the initial amount of wealth.

A decision maker with diminishing marginal utility for wealth will be risk averse.

Much more likely to take her chances, as others do when faced with very bad options.

Once you have accepted a theory and used it as a tool in your thinking, it is extraordinarily difficult to notice its flaws. If you come upon an observation that does not seem to fit the model, you assume that there must be a perfectly good explanation that you are somehow missing. You give the theory the benefit of the doubt, trusting the community of experts who have accepted it.

Disbelieving is hard work, and System 2 is easily tired.

People become risk seeking when all their options are bad.

Organisms that treat threats as more urgent than opportunities have a better chance to survive and reproduce.

In mixed gambles, where both a gain and a loss are possible, loss aversion causes extremely risk-averse choices. In bad choices, where a sure loss is compared to a larger loss that is merely probable, diminishing sensitivity causes risk seeking.

Diminishing marginal utility: the more leisure you have, the less you care for an extra day of it, and each added day is worth less than the one before. Similarly, the more income you have, the less you care for an extra dollar, and the amount you are willing to give up for an extra day of leisure increases.

A mistaken assumption: that your utility for a state of affairs depends only on that state and is not affected by your history.

When you shop for shoes, the merchant who gives up the shoes in exchange for money certainly feels no loss. Indeed, the shoes that he hands over have always been, from his point of view, a cumbersome proxy for money that he was hoping to collect from some consumer. Furthermore, you probably do not experience paying the merchant as a loss, because you were effectively holding money as a proxy for the shoes you intended to buy.

Both the shoes the merchant sells you and the money you spend from your budget for shoes are held “for exchange.” They are intended to be traded for other goods. Other goods, such as wine and Super Bowl tickets, are held “for use,” to be consumed or otherwise enjoyed.

Your leisure time and the standard of living that your income supports are also not intended for sale or exchange.

Only 18% of the inexperienced traders were willing to exchange their gift for the other. In sharp contrast, experienced traders showed no trace of an endowment effect: 48% of them traded!

People who are poor think like traders, but all their choices are between losses. Money that is spent on one good is the loss of another good that could have been purchased instead. For the poor, costs are losses.

Asked a friend whether he would accept a gamble on the toss of a coin in which he could lose \$100 or win \$200. His friend responded, "I won't bet because I would feel the \$100 loss more than the \$200 gain. But I'll take you on if you promise to let me make 100 such bets."

I sympathize with your aversion to losing any gamble, but it is costing you a lot of money.

Is this the last offer of a small favorable gamble that you will ever consider?

You will have many opportunities to consider attractive gambles with stakes that are very small relative to your wealth. You will do yourself a large financial favor if you are able to see each of these gambles as part of a bundle of small gambles.

Rehearse the mantra that will get you significantly closer to economic rationality: you win a few, you lose a few. The main purpose of the mantra is to control your emotional response when you do lose.

The mantra works when the gambles are genuinely independent of each other; it does not apply to multiple investments in the same industry, which would all go bad together. It works only when the possible loss does not cause you to worry about your total wealth. If you would take the loss as significant bad news about your economic future, watch it! It should not be applied to long shots, where the probability of winning is very small for each bet. If you have the emotional discipline that this rule requires, you will never consider a small gamble in isolation or be loss averse for a small gamble.

Broad framing blunted the emotional reaction to losses and increased the willingness to take risks. The combination of loss aversion and narrow framing is a costly curse. Individual investors can avoid that curse, achieving the emotional benefits of broad framing while also saving time and agony, by reducing the frequency with which they check how well their investments are doing. Closely following daily fluctuations is a losing proposition, because the pain of the frequent small losses exceeds the pleasure of the equally frequent small gains.

The deliberate avoidance of exposure to short-term outcomes improves the quality of both decisions and outcomes.

A commitment not to change one's position for several periods (the equivalent of "locking in" an investment) improves financial performance.

Having a risk policy that they routinely apply whenever a relevant problem arises. Familiar examples of risk policies are "always take the highest possible deductible when purchasing insurance" and "never buy extended warranties." A risk policy is a broad frame.

Reduce or eliminate the pain of the occasional loss by the thought that the policy that left you exposed to it will almost certainly be financially advantageous over the long run.

The outside view and the risk policy are remedies against two distinct biases that affect many decisions: the exaggerated optimism of the planning fallacy and the exaggerated caution induced by loss aversion.

Top managers of the 25 divisions of a large company. He asked them to consider a risky option in which, with equal probabilities, they could lose a large amount of the capital they controlled or earn double that amount. None of the executives was willing to take such a dangerous gamble. Thaler then turned to the CEO of the company, who was also present, and asked for his opinion. Without hesitation,

the CEO answered, “I would like all of them to accept their risks.” In the context of that conversation, it was natural for the CEO to adopt a broad frame that encompassed all 25 bets.

He could count on statistical aggregation to mitigate the overall risk.

Money is a proxy for points on a scale of self-regard and achievement.

Finance research has documented a massive preference for selling winners rather than losers - a bias that has been given an opaque label: the disposition effect. The disposition effect is an instance of narrow framing. The investor has set up an account for each share that she bought, and she wants to close every account as a gain. A rational agent would have a comprehensive view of the portfolio.

The sunk-cost fallacy keeps people for too long in poor jobs, unhappy marriages, and unpromising research projects.

=====

ALGORITHMS

Each of these domains entails a significant degree of uncertainty and unpredictability. We describe them as “low-validity environments.” In every case, the accuracy of experts was matched or exceeded by a simple algorithm.

Clinical vs. Statistical Prediction: A Theoretical Analysis and a Review of the Evidence.

Orley Ashenfelter has offered a compelling demonstration of the power of simple statistics to outdo world-renowned experts. Ashenfelter wanted to predict the future value of fine Bordeaux wines from information available in the year they are made.

Ashenfelter converted that conventional knowledge into a statistical formula that predicts the price of a wine - for a particular property and at a particular age - by three features of the weather: the

average temperature over the summer growing season, the amount of rain at harvest-time, and the total rainfall during the previous winter. His formula provides accurate price forecasts years and even decades into the future. Indeed, his formula forecasts future prices much more accurately than the current prices of young wines do.

Ashenfelter's formula is extremely accurate - the correlation between his predictions and actual prices is above .90.

Why are experts inferior to algorithms? One reason, which Meehl suspected, is that experts try to be clever, think outside the box, and consider complex combinations of features in making their predictions. Complexity may work in the odd case, but more often than not it reduces validity. Simple combinations of features are better.

Human decision makers are inferior to a prediction formula even when they are given the score suggested by the formula! They feel that they can overrule the formula because they have additional information.

There are few circumstances under which it is a good idea to substitute judgment for a formula. In a famous thought experiment, he described a formula that predicts whether a particular person will go to the movies tonight and noted that it is proper to disregard the formula if information is received that the individual broke a leg today. The name "broken-leg rule" has stuck. The point, of course, is that broken legs are very rare - as well as decisive.

To maximize predictive accuracy, final decisions should be left to formulas, especially in low-validity environments.

It is possible to develop useful algorithms without any prior statistical research. Simple equally weighted formulas based on existing statistics or on common sense are often very good predictors of significant outcomes.

Marital stability is well predicted by a formula: frequency of lovemaking minus frequency of quarrels.

An algorithm that is constructed on the back of an envelope is often good enough to compete with an optimally weighted formula, and certainly good enough to outdo expert judgment.

If you are serious about hiring the best possible person for the job, this is what you should do. First, select a few traits that are prerequisites for success in this position (technical proficiency, engaging personality, reliability, and so on). Don't overdo it - six dimensions is a good number. The traits you choose should be as independent as possible from each other, and you should feel that you can assess them reliably by asking a few factual questions. Next, make a list of those questions for each trait and think about how you will score it, say on a 1–5 scale.

Collect the information on one trait at a time, scoring each before you move on to the next one. Do not skip around. To evaluate each candidate, add up the six scores.

Firmly resolve that you will hire the candidate whose final score is the highest, even if there is another one whom you like better - try to resist your wish to invent broken legs to change the ranking.

=====
OUTSIDE VIEW

Baseline prediction: the prediction you make about a case if you know nothing except the category to which it belongs.

The baseline prediction should be the anchor for further adjustments.

People who have information about an individual case rarely feel the need to know the statistics of the class to which the case belongs.

“What is the probability of the defendant winning in cases like this one?” His sharp answer: “Every case is unique.”

A proud emphasis on the uniqueness of cases is also common in medicine, in spite of recent advances in evidence-based medicine that point the other way.

A survey of American homeowners who had remodeled their kitchens found that, on average, they had expected the job to cost \$18,658; in fact, they ended up paying an average of \$38,769.

The greatest responsibility for avoiding the planning fallacy lies with the decision makers who approve the plan.

If they do not recognize the need for an outside view, they commit a planning fallacy.

The prevalent tendency to underweight or ignore distributional information is perhaps the major source of error in forecasting. Planners should therefore make every effort to frame the forecasting problem so as to facilitate utilizing all the distributional information that is available. This may be considered the single most important piece of advice regarding how to increase accuracy in forecasting through improved methods.

Identify an appropriate reference class (kitchen renovations, large railway projects, etc.). Obtain the statistics of the reference class (in terms of cost per mile of railway, or of the percentage by which expenditures exceeded budget). Use the statistics to generate a baseline prediction. Use specific information about the case to adjust the baseline prediction, if there are particular reasons to expect the optimistic bias to be more or less pronounced in this project than in others of the same type.

Executives too easily fall victim to the planning fallacy. In its grip, they make decisions based on delusional optimism rather than on a rational weighting of gains, losses, and probabilities. They

overestimate benefits and underestimate costs. They spin scenarios of success while overlooking the potential for mistakes and miscalculations. As a result, they pursue initiatives that are unlikely to come in on budget or on time or to deliver the expected returns - or even to be completed. In this view, people often (but not always) take on risky projects because they are overly optimistic about the odds they face.

This is an explanation of why people litigate, why they start wars, and why they open small businesses.

If you are genetically endowed with an optimistic bias, you hardly need to be told that you are a lucky person - you already feel fortunate. An optimistic attitude is largely inherited, and it is part of a general disposition for well-being, which may also include a preference for seeing the bright side of everything. If you were allowed one wish for your child, seriously consider wishing him or her optimism. Optimists are normally cheerful and happy, and therefore popular; they are resilient in adapting to failures and hardships, their chances of clinical depression are reduced, their immune system is stronger, they take better care of their health, they feel healthier than others and are in fact likely to live longer. A study of people who exaggerate their expected life span beyond actuarial predictions showed that they work longer hours, are more optimistic about their future

Optimistic individuals play a disproportionate role in shaping our lives. Their decisions make a difference; they are the inventors, the entrepreneurs, the political and military leaders - not average people. They got to where they are by seeking challenges and taking risks. They are talented and they have been lucky, almost certainly luckier than they acknowledge.

The people who have the greatest influence on the lives of others are likely to be optimistic and overconfident, and to take more risks than they realize.

These persistent (or obstinate) individuals doubled their initial losses before giving up. Significantly, persistence after discouraging advice was relatively common among inventors who had a high score on a personality measure of optimism.

The damage caused by overconfident CEOs is compounded when the business press anoints them as celebrities; the evidence indicates that prestigious press awards to the CEO are costly to stockholders. The authors write, "We find that firms with award-winning CEOs subsequently underperform, in terms both of stock and of operating performance. At the same time, CEO compensation increases, CEOs spend more time on activities outside the company such as writing books and sitting on outside boards, and they are more likely to engage in earnings management."

To explain entrepreneurial optimism, cognitive biases play an important role.

We focus on our goal, anchor on our plan, and neglect relevant base rates, exposing ourselves to the planning fallacy. We focus on what we want to do and can do, neglecting the plans and skills of others. Both in explaining the past and in predicting the future, we focus on the causal role of skill and neglect the role of luck. We are therefore prone to an illusion of control. We focus on what we know and neglect what we do not know, which makes us overly confident in our beliefs.

I have had several occasions to ask founders and participants in innovative start-ups a question: To what extent will the outcome of your effort depend on what you do in your firm? This is evidently an easy question; the answer comes quickly and in my small sample it has never been less than 80%. Even when they are not sure they will succeed, these bold people think their fate is almost entirely in their own hands. They are surely wrong: the outcome of a start-up depends as much on the achievements of its competitors and on changes in the market as on its own efforts. However, WY SIATI plays

its part, and entrepreneurs naturally focus on what they know best - their plans and actions and the most immediate threats and opportunities, such as the availability of funding. They know less about their competitors and therefore find it natural to imagine a future in which the competition plays little part.

Entrepreneurial firms that fail but signal new markets to more qualified competitors “optimistic martyrs” - good for the economy but bad for their investors.

A survey in which the chief financial officers of large corporations estimated the returns of the Standard & Poor's index over the following year. The Duke scholars collected 11,600 such forecasts and examined their accuracy. The conclusion was straightforward: financial officers of large corporations had no clue about the short-term future of the stock market; the correlation between their estimates and the true value was slightly less than zero! When they said the market would go down, it was slightly more likely than not that it would go up.

The answer that a truthful CFO would offer is plainly ridiculous. A CFO who informs his colleagues that “there is a good chance that the S&P returns will be between -10% and +30%” can expect to be laughed out of the room. The wide confidence interval is a confession of ignorance, which is not socially acceptable for someone who is paid to be knowledgeable in financial matters. Even if they knew how little they know, the executives would be penalized for admitting it.

The emotional, cognitive, and social factors that support exaggerated optimism are a heady brew, which sometimes leads people to take risks that they would avoid if they knew the odds.

The contribution of optimism to good implementation is certainly positive. The main benefit of optimism is resilience in the face of setbacks.

Someone who lacks a delusional sense of significance will wilt in the face of repeated experiences of multiple small failures and rare successes, the fate of most small business.

When the organization has almost come to an important decision but has not formally committed itself, Klein proposes gathering for a brief session a group of individuals who are knowledgeable about the decision. The premise of the session is a short speech: "Imagine that we are a year into the future. We implemented the plan as it now exists. The outcome was a disaster. Please take 5 to 10 minutes to write a brief history of that disaster."

Premortem has two main advantages: it overcomes the groupthink that affects many teams once a decision appears to have been made, and it unleashes the imagination of knowledgeable individuals in a much-needed direction.

=====

LOSSES

We submitted our essay to *Econometrica*, a journal that publishes significant theoretical articles in economics and in decision theory. The choice of venue turned out to be important; if we had published the identical paper in a psychological journal, it would likely have had little impact on economics. However, our decision was not guided by a wish to influence economics; *Econometrica* just happened to be where the best papers on decision making had been published in the past, and we were aspiring to be in that company.

A single cockroach will completely wreck the appeal of a bowl of cherries, but a cherry will do nothing at all for a bowl of cockroaches.

Bad emotions, bad parents, and bad feedback have more impact than good ones, and bad information is processed more thoroughly than good. The self is more motivated to avoid bad self-definitions than to pursue good ones. Bad impressions and bad stereotypes are

quicker to form and more resistant to disconfirmation than good ones.

The long-term success of a relationship depends far more on avoiding the negative than on seeking the positive.

A friendship that may take years to develop can be ruined by a single action.

The aversion to the failure of not reaching the goal is much stronger than the desire to exceed it. People often adopt short-term goals that they strive to achieve but not necessarily to exceed. They are likely to reduce their efforts when they have reached an immediate goal, with results that sometimes violate economic logic.

Players would try a little harder when putting for par (to avoid a bogey) than when putting for a birdie.

Loss aversion creates an asymmetry that makes agreements difficult to reach. The concessions you make to me are my gains, but they are your losses; they cause you much more pain than they give me pleasure. Inevitably, you will place a higher value on them than I do.

Negotiators often pretend intense attachment to some good.

Although they actually view that good as a bargaining chip and intend ultimately to give it away in an exchange. Because negotiators are influenced by a norm of reciprocity, a concession that is presented as painful calls for an equally painful (and perhaps equally inauthentic) concession from the other side.

A biologist observed that “when a territory holder is challenged by a rival, the owner almost always wins the contest - usually within a matter of seconds.”

Altruistic punishment is accompanied by increased activity in the “pleasure centers” of the brain. It appears that maintaining the social order and the rules of fairness in this fashion is its own reward.

Altruistic punishment could well be the glue that holds societies together.

People attach values to gains and losses rather than to wealth.

The fourfold pattern of preferences is considered one of the core achievements of prospect theory.

When you consider a choice between a sure loss and a gamble with a high probability of a larger loss, diminishing sensitivity makes the sure loss more aversive, and the certainty effect reduces the aversiveness of the gamble.

This is where people who face very bad options take desperate gambles, accepting a high probability of making things worse in exchange for a small hope of avoiding a large loss. Risk taking of this kind often turns manageable failures into disasters. The thought of accepting the large sure loss is too painful, and the hope of complete relief too enticing, to make the sensible decision that it is time to cut one's losses. This is where businesses that are losing ground to a superior technology waste their remaining assets in futile attempts to catch up. Because defeat is so difficult to accept, the losing side in wars often fights long past the point at which the victory of the other side is certain.

People expect to have stronger emotional reactions (including regret) to an outcome that is produced by action than to the same outcome when it is produced by inaction.

Be explicit about the anticipation of regret. People generally anticipate more regret than they will actually experience, because they underestimate the efficacy of the psychological defenses they will deploy - which they label the "psychological immune system." Their recommendation is that you should not put too much weight on regret; even if you have some, it will hurt less than you now think.

When you see cases in isolation, you are likely to be guided by an emotional reaction of System 1.

Would you accept a gamble that offers a 10% chance to win \$95 and a 90% chance to lose \$5? Would you pay \$5 to participate in a lottery that offers a 10% chance to win \$100 and a 90% chance to win nothing?

A bad outcome is much more acceptable if it is framed as the cost of a lottery ticket that did not win than if it is simply described as losing a gamble. We should not be surprised: losses evokes stronger negative feelings than costs.

Tendencies to approach or avoid are evoked by the words, and we expect System 1 to be biased in favor of the sure option when it is designated as KEEP and against that same option when it is designated as LOSE.

Decision makers tend to prefer the sure thing over the gamble (they are risk averse) when the outcomes are good. They tend to reject the sure thing and accept the gamble (they are risk seeking) when both outcomes are negative.

System 1 delivers an immediate response to any question about rich and poor: when in doubt, favor the poor.

Your moral feelings are attached to frames, to descriptions of reality rather than to reality itself.

=====
EXPERIENCING

The experiencing self is the one that answers the question: "Does it hurt now?" The remembering self is the one that answers the question: "How was it, on the whole?" Memories are all we get to keep from our experience of living, and the only perspective that we

can adopt as we think about our lives is therefore that of the remembering self.

Confusing experience with the memory of it is a compelling cognitive illusion - and it is the substitution that makes us believe a past experience can be ruined. The experiencing self does not have a voice. The remembering self is sometimes wrong, but it is the one that keeps score and governs what we learn from living, and it is the one that makes decisions. What we learn from the past is to maximize the qualities of our future memories, not necessarily of our future experience.

My happy childhood: I always cried when my mother came to tear me away from my toys to take me to the park, and cried again when she took me away from the swings and the slide. The resistance to interruption was a sign I had been having a good time, both with my toys and with the swings.

At random intervals during the day, the phone presents a brief menu of questions about what the respondent was doing and who was with her when she was interrupted. The participant is also shown rating scales to report the intensity of various feelings: happiness, tension, anger, worry, engagement, physical pain, and others.

The percentage of time that an individual spends in an unpleasant state is the U-index. For example, an individual who spent 4 hours of a 16-hour waking day in an unpleasant state would have a U-index of 25%.

Our emotional state is largely determined by what we attend to, and we are normally focused on our current activity and immediate environment.

To get pleasure from eating, for example, you must notice that you are doing it.

Americans were far more prone to combine eating with other activities, and their pleasure from eating was correspondingly diluted.

Priming students with the idea of wealth reduces the pleasure their face expresses as they eat a bar of chocolate.

Goals make a large difference. Nineteen years after they stated their financial aspirations, many of the people who wanted a high income had achieved it. Each additional point on the money-importance scale was associated with an increment of over \$14,000 of job income.

The goals that people set for themselves are so important to what they do and how they feel about it that an exclusive focus on experienced well-being is not tenable. We cannot hold a concept of well-being that ignores what people want.

Nothing in life is as important as you think it is when you are thinking about it.

Source: [Thinking, Fast and Slow - by Daniel Kahneman | Derek Sivers](#)