INTELLIGENCE · ANALYSIS AND TRADECRAFT

Cognitive Biases in Intelligence Analysis

Analysts, like everyone, rely on mental shortcuts that can lead to systematic errors. This article explains the cognitive biases that matter most in intelligence work, how they have contributed to real failures, and which countermeasures research suggests actually help.

ModQu research article · 2,248 words · 22 references (APA 7th edition)

Intelligence analysis depends on human judgement. Analysts select information, interpret it, weigh alternatives and estimate what is likely to happen. Human judgement is remarkably capable, but it is not neutral. Decades of research in psychology have shown that people rely on mental shortcuts, called **heuristics**, that usually work well but produce predictable errors, called **cognitive biases**. Because the errors are systematic rather than random, they do not cancel out; they push judgements consistently in the wrong direction. This article explains the biases that matter most for intelligence work, how they have contributed to failures, and what can be done about them.

Where the research comes from

The modern study of judgement under uncertainty began with the work of Amos Tversky and Daniel Kahneman. In an influential paper, Tversky and Kahneman (1974) described three heuristics, representativeness, availability, and anchoring and adjustment, and the biases each produced. Kahneman (2011) later summarised this research for a general audience, describing two modes of thinking: a fast, automatic and intuitive mode, and a slow, deliberate and effortful one. Many biases arise when the fast mode answers a question that needs the slow one.

Richards Heuer brought this research to intelligence analysts. In a book on the psychology of intelligence analysis, Heuer (1999) explained how perception, memory and reasoning shape the way analysts handle evidence, and argued that simply knowing about biases does little to prevent them. Heuer's work underpins much of the structured analytic tradition (Pherson & Heuer, 2021).

The main biases

Common biases in analysis and how to counter them

Each bias has a characteristic effect on judgement and a practical countermeasure

Bias	Effect on analysis	Countermeasure
Confirmation	Evidence for a favoured view is sought and believed	List alternatives; seek disconfirming evidence
Anchoring	Estimates stay close to a first figure	Make independent estimates; consider the opposite
Availability	Vivid or recent events judged more likely	Check base rates and records
Base-rate neglect	Specific details outweigh how common things are	Start from the reference class
Overconfidence	Ranges too narrow; certainty too high	State probabilities; track accuracy
Mirror imaging	Others assumed to think as we do	Red team; consult people who know the actor
Hindsight	Past events seem to have been predictable	Record judgements before outcomes are known
Conservatism	Views updated too slowly on new evidence	Set indicators; review on a schedule

Compiled from Tversky and Kahneman (1974), Heuer (1999) and Nickerson (1998)

ModQu.com

Figure 1. Common biases in analysis and how to counter them. Compiled from Tversky and Kahneman (1974), Heuer (1999) and Nickerson (1998); table by ModQu.

Confirmation bias

Confirmation bias is the tendency to look for, interpret and remember information in ways that support what one already believes. Nickerson (1998) reviewed evidence of its many forms. Lord et al. (1979) showed that people presented with the same mixed evidence on a contested issue each found it supported their original view, and became more confident as a result. Klayman and Ha (1987) explained that a natural "positive test strategy", looking for cases where a hypothesis should hold, is often reasonable but fails when confirming evidence would also fit other explanations.

In intelligence, confirmation bias leads analysts to accumulate supporting evidence for a favoured judgement while discounting or explaining away contrary evidence. It is perhaps the most important single bias, because it affects how all other evidence is processed.

Anchoring

When people estimate a quantity, they tend to start from an initial value and adjust insufficiently. Tversky and Kahneman (1974) showed that even obviously arbitrary starting numbers influenced estimates. In analysis, an initial estimate, a previous assessment or a figure in a first report can act as an anchor that later evidence fails to move far enough.

Availability

Tversky and Kahneman (1973) showed that people judge the frequency or likelihood of events by how easily examples come to mind. Vivid, recent or widely reported events seem more probable than they are; quiet, gradual or unreported developments seem less so. An analyst who has just handled a dramatic cyber attack may overestimate the likelihood of another; one who has never seen a particular kind of event may dismiss it.

Representativeness and base-rate neglect

People judge how likely something is by how closely it resembles a typical case. Bar-Hillel (1980) reviewed evidence that this leads people to neglect **base rates**, the background frequency of events. If an analyst is told that a company shows several signs typical of firms about to be acquired, the resemblance may seem compelling, even though only a small fraction of firms are acquired in any year. Kahneman and Lovallo (1993) described the remedy as taking the **outside view**: starting from what usually happens in similar cases before adjusting for the specifics.

Overconfidence

Moore and Healy (2008) distinguish three forms of overconfidence: overestimating one's own ability or performance, believing oneself better than others, and being too sure that one's estimates are correct. The third, **overprecision**, is especially relevant to analysis. People asked to give a range that they are 90 per cent sure contains the true value typically give ranges that contain it far less often.

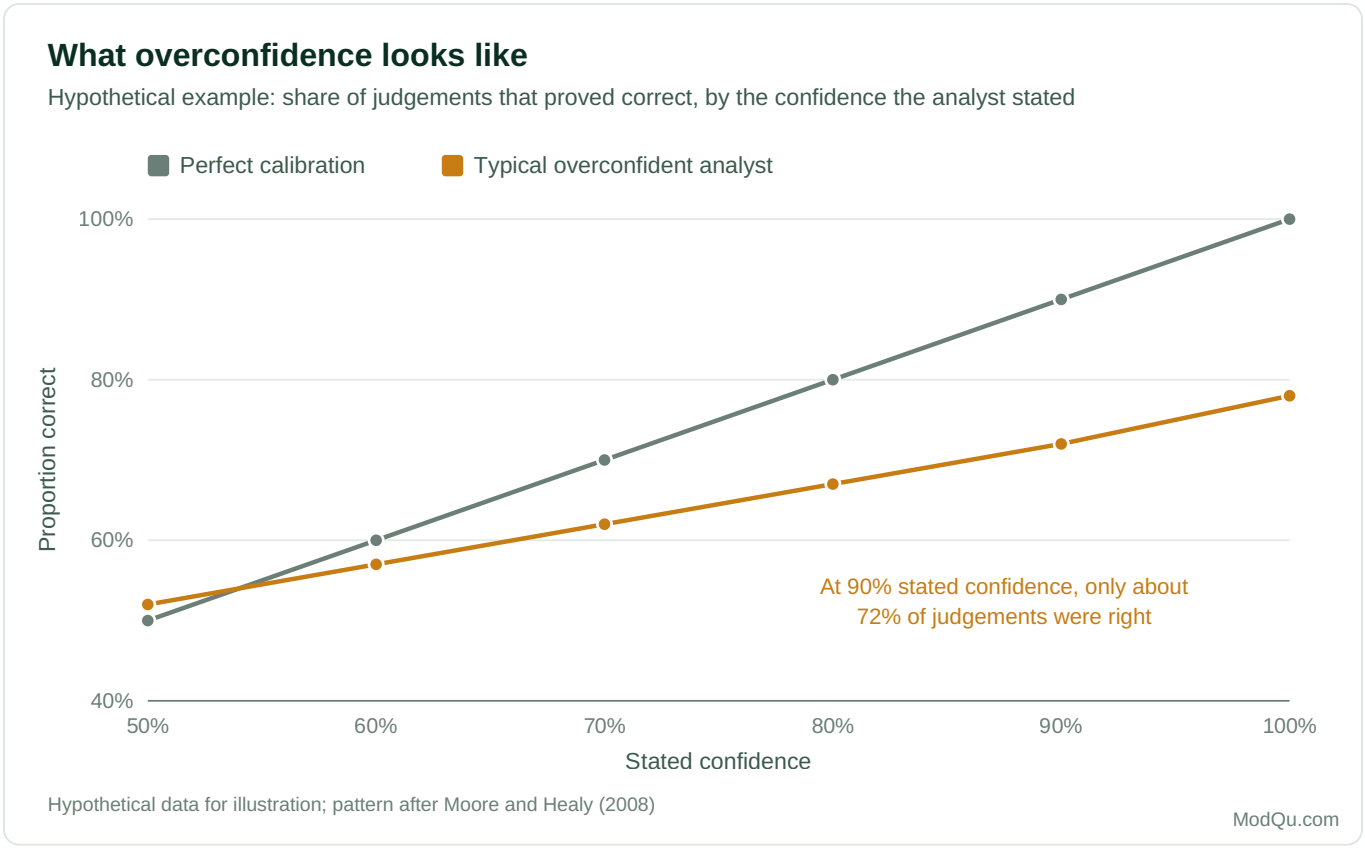


Figure 2. What overconfidence looks like. When stated confidence is compared with actual accuracy, overconfident judgements fall below the diagonal. Hypothetical data for illustration; pattern after Moore and Healy (2008); chart by ModQu.

Koriat et al. (1980) found that asking people to list reasons why their answer might be wrong reduced overconfidence, while listing reasons in favour did not.

Hindsight bias

Fischhoff (1975) showed that once people know an outcome, they overestimate how predictable it was and misremember what they had believed beforehand. Hindsight bias distorts the evaluation of intelligence: after a surprise, the warning signs look obvious, and critics conclude that analysts should have seen them. It also distorts learning, because analysts themselves may believe they had expected an outcome that in fact surprised them.

Conservatism

Phillips and Edwards (1966) found that people revise their beliefs in response to new evidence less than probability theory says they should. In analysis, this appears as a reluctance to abandon an established view, so that accumulating evidence of change is absorbed bit by bit without altering the overall judgement.

Framing

Tversky and Kahneman (1981) showed that logically equivalent descriptions of the same choice, for example in terms of lives saved or lives lost, lead to different decisions. How an intelligence question is framed, and how a judgement is presented, can shape both the analysis and the response of decision-makers.

Mirror imaging

Heuer (1999) described **mirror imaging** as the assumption that others will think and act as the analyst would in their situation. It is a particular danger when analysing foreign governments, unfamiliar cultures or competitors with different incentives. A competitor may accept losses that the analyst's own firm would never tolerate, because its owners have different goals.

Groupthink

Some biases arise in groups rather than individuals. Janis (1972), studying American foreign policy decisions, described **groupthink**: the tendency of cohesive groups under pressure to seek agreement, suppress doubts and fail to consider alternatives.

Biases in real failures

Jervis (2010) analysed two major intelligence failures in detail: the failure to anticipate the fall of the Shah of Iran in 1979, and the mistaken judgement that Iraq possessed weapons of mass destruction before the 2003 war. In both cases Jervis found that analysts' reasoning was shaped by strong prior beliefs that were rarely examined. In the Iraq case, Iraq's past deception and evasive behaviour were read as evidence of hidden weapons, when other explanations, such as a wish to appear stronger to regional rivals, were also consistent with what was observed. Contrary evidence was explained away; alternative hypotheses were not seriously tested. Jervis also cautions against hindsight: some of the evidence genuinely pointed in the wrong direction, and the failures were not simply the result of carelessness.

Noise as well as bias

Kahneman et al. (2021) emphasise a second source of error that receives less attention: **noise**, meaning unwanted variability in judgements that should be the same. Different analysts given the same evidence may reach very different conclusions, and the same analyst may judge differently on different days. Bias shifts judgements in one direction; noise scatters them. Both reduce accuracy, and noise can be reduced through structured procedures, independent judgements combined by averaging, and clear scales for expressing conclusions.

Biases in business intelligence

The same biases appear in commercial settings, often in familiar forms:

- **Planning and forecasting.** Kahneman and Lovallo (1993) describe how managers estimating the cost and timing of new projects tend to take the "inside view", focusing on the details of their own plan, and produce forecasts that are consistently too optimistic. Comparing the plan with the record of similar projects corrects much of the error.
- **Competitor analysis.** Mirror imaging leads firms to assume competitors share their goals and constraints. A rival backed by owners who value growth over profit, or a state-supported entrant, may act in ways the analyst's own firm would consider irrational.
- **Market sizing.** Anchoring on a figure from an early report, or on a senior executive's estimate, can shape every later calculation.
- **Customer research.** Confirmation bias leads teams to read survey results and customer comments as support for a product idea they already favour, while discounting signs of weak demand.
- **Post-mortems.** Hindsight bias makes failed launches or missed threats look obvious afterwards, which can lead to blame rather than learning.

Pressures that amplify bias

Biases are strongest when analysts work under conditions that favour fast, intuitive judgement: tight deadlines, heavy workloads, ambiguous evidence and pressure to give a clear answer. Kahneman (2011) notes that the deliberate mode of thinking is effortful and easily depleted, so tired or rushed analysts fall back on shortcuts. Organisational incentives also matter. If an analyst's standing depends on being seen as decisive, or if contradicting a senior colleague carries a cost, the social pressure to agree reinforces confirmation bias and conservatism. Recognising these conditions allows managers to schedule the most important judgements for when analysts have time to think, and to build in a pause for structured challenge before major conclusions are issued.

What helps

Knowing is not enough

Pronin et al. (2002) found that people readily recognise biases in others but believe themselves to be less affected, a pattern called the **bias blind spot**. Heuer (1999) made the same point: awareness of a bias does not by itself remove it. Effective countermeasures change the procedure, not just the mindset.

Consider the opposite

Among the best-supported techniques is to **consider the opposite**. Lord et al. (1984) found that asking people to consider how they would have judged the evidence if it had pointed the other way was more effective at reducing bias than simply instructing them to be fair. Mussweiler et al. (2000) showed that considering reasons why an anchor was inappropriate reduced its effect. In analysis, the equivalent is to ask, for any important judgement, what the evidence would look like if it were wrong, and to look for that evidence.

Use base rates and reference classes

Starting with the base rate, how often events of this kind occur in comparable cases, and then adjusting for specific evidence, counters both representativeness and overconfidence. Mandel (2015) found that brief instruction in structuring information as frequencies improved the accuracy of intelligence analysts' probability judgements.

State probabilities and track accuracy

Expressing judgements as probabilities makes overconfidence visible, and checking those probabilities against outcomes over time shows each analyst and team where they are miscalibrated.

Structure and challenge

Listing assumptions, setting out alternative hypotheses, inviting dissent and using red teams all counter the tendency to settle on a single favoured view. Pherson and Heuer (2021) describe many such techniques, although, as discussed elsewhere, the experimental evidence for individual techniques is mixed.

Record judgements in advance

Writing down judgements, with probabilities and reasoning, before outcomes are known protects against hindsight bias and makes honest learning possible. A simple log, recording the date, the question, the judgement, the probability and the main evidence, is enough. Reviewed every few months, it shows which kinds of questions a team handles well and where its confidence runs ahead of its accuracy.

Organisational measures

Some of the most effective defences operate at the level of the organisation rather than the individual:

1. **Separate independent judgements.** Ask analysts to form views independently before discussion, so that the first speaker or the most senior person does not anchor the group.
2. **Make dissent safe.** Reward people who raise well-reasoned objections, even when they turn out to be wrong.
3. **Review products for reasoning, not just conclusions.** Editors and managers should ask how alternatives were considered and what would change the judgement.
4. **Learn from outcomes.** Compare past judgements with what happened, without blame, and share the lessons.
5. **Bring in outsiders.** People from other teams, disciplines or backgrounds see different possibilities.

An illustration

A manufacturer's intelligence team has assessed for three years that a foreign competitor lacks the technology to enter its market. Each new item of information, a trade show appearance, a hiring announcement, a patent filing, is noted and judged "not significant". When a new analyst joins and is asked to consider the opposite, the newcomer lists what the team would expect to see if the competitor were preparing to enter: patents in the relevant technology, recruitment of engineers with specific skills, certification applications and visits to potential distributors. The newcomer finds that most of these have occurred in the past eighteen months. Each had been explained away individually, a pattern of conservatism and confirmation bias. The team revises its assessment, assigns a probability of entry within two years and sets indicators to monitor. It also changes its routine: new information about the competitor is now logged against a short list of hypotheses, so that a pattern of small changes is visible rather than being judged one item at a time, and the assessment is formally reviewed every quarter by someone who did not write it.

Conclusion

Cognitive biases are not signs of carelessness or poor training; they are features of how human minds work. In intelligence analysis, where evidence is incomplete and stakes are high, they can lead to serious errors, as the history of intelligence failures shows. Awareness alone does little. What helps is changing procedures: considering the opposite, starting from base rates, stating probabilities and checking them, structuring analysis to examine alternatives, making dissent safe and recording judgements so that the organisation can learn honestly from experience.

References

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, *44*(3), 211–233. https://doi.org/10.1016/0001-6918(80)90046-3

Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance, 1*(3), 288–299. https://doi.org/10.1037/0096-1523.1.3.288

Heuer, R. J., Jr. (1999). *Psychology of intelligence analysis*. Center for the Study of Intelligence, Central Intelligence Agency.

Janis, I. L. (1972). *Victims of groupthink: A psychological study of foreign-policy decisions and fiascoes*. Houghton Mifflin.

Jervis, R. (2010). *Why intelligence fails: Lessons from the Iranian Revolution and the Iraq War*. Cornell University Press.

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science, 39*(1), 17–31. https://doi.org/10.1287/mnsc.39.1.17

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. Little, Brown Spark.

Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94*(2), 211–228. https://doi.org/10.1037/0033-295X.94.2.211

Koriat, A., Lichtenstein, S., & Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory, 6*(2), 107–118. https://doi.org/10.1037/0278-7393.6.2.107

Lord, C. G., Lepper, M. R., & Preston, E. (1984). Considering the opposite: A corrective strategy for social judgment. *Journal of Personality and Social Psychology, 47*(6), 1231–1243. https://doi.org/10.1037/0022-3514.47.6.1231

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*(11), 2098–2109. https://doi.org/10.1037/0022-3514.37.11.2098

Mandel, D. R. (2015). Instruction in information structuring improves Bayesian judgment in intelligence analysts. *Frontiers in Psychology, 6*, Article 387. https://doi.org/10.3389/fpsyg.2015.00387

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review, 115*(2), 502–517. https://doi.org/10.1037/0033-295X.115.2.502

Mussweiler, T., Strack, F., & Pfeiffer, T. (2000). Overcoming the inevitable anchoring effect: Considering the opposite compensates for selective accessibility. *Personality and Social Psychology Bulletin, 26*(9), 1142–1150. https://doi.org/10.1177/01461672002611010

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175–220. https://doi.org/10.1037/1089-2680.2.2.175

Pherson, R. H., & Heuer, R. J., Jr. (2021). *Structured analytic techniques for intelligence analysis* (3rd ed.). CQ Press.

Phillips, L. D., & Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology, 72*(3), 346–354. https://doi.org/10.1037/h0023653

Pronin, E., Lin, D. Y., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin, 28*(3), 369–381. https://doi.org/10.1177/0146167202286008

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5*(2), 207–232. https://doi.org/10.1016/0010-0285(73)90033-9

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*(4481), 453–458. https://doi.org/10.1126/science.7455683