



WHITE PAPER • OCTOBER 2026

What bias does to a decision

Twelve habits that shape executive judgment, and the moment each one can be caught

Overtly Learning

Written by Orlando Franco



Every experienced decision maker runs on the same handful of shortcuts. Knowing which one is working on you is the part you can act on.

Bias is not a failing of careless people. It is what judgment does when information is incomplete and time is short, which describes every decision worth the name. What follows is twelve of them, what each one does to a decision, and the moment it can be caught.

about a third

of answers bend to a confident majority giving an obviously wrong answer

2 to 4 in 10

how often ranges meant to be right nine times in ten actually caught the answer


8 of 8

categories of capital project where average cost came in above the estimate

- **Nobody is outside it.** People rate themselves as less prone to bias than the people around them, and that finding has survived careful retesting. When somebody looked for a link with cognitive ability, they did not find one.
- **The first family acts on your numbers.** Where an estimate starts, how an option is worded, how wide a range is, and how a schedule is built. These are the quiet ones, because the output still looks like arithmetic.
- **The second is about being right rather than getting it right.** Seeking the evidence that agrees, claiming wins and explaining away losses, and spending more because you have already spent.
- **The third is what the room does** to a judgment that was sound when it walked in: who speaks first, what everyone already knows, and what nobody says.
- **What helps changes the process, not the person.** The countermeasures aimed at making individuals less biased have thin records. The ones that change how a decision is made have better ones.

THE ASK

Pick one recurring decision and change how it is made, not who makes it.



1 · NOBODY IS OUTSIDE THIS

Ask a leadership team who is most prone to bias. Few people point at themselves.

People rate themselves as less prone to bias than the people around them, and the gap is wide. In the original study, most of the people who had rated themselves above average did not revise that rating after reading a description of the very bias they had just displayed. They insisted it was accurate, or even modest.

Being clever does not protect you. An earlier study went looking for a link with cognitive ability and did not find one, and the biases themselves appear to work outside awareness and outside choice.

That is why a paper like this one is worth more than a self-assessment. You cannot inspect your own judgment for these; you can learn what they look like from outside and change the conditions that let them operate.

One honest caveat before the list

This field spent years retesting its own famous results. About half held up, and three in four came back smaller than first reported. Replicated usually means smaller, not confirmed. So how often you have heard a claim about decision making, and how confidently it is repeated, tells you very little about whether it will hold.

Everything in this paper is drawn from a course corpus of open, published research, and where something is widely recommended but has never actually been tested, the paper says so rather than leaving you to assume.

EXHIBIT 1

Twelve habits, three families, and the moment each one shows up

Everyone uses these. The useful question is never whether you are biased, it is which of these is operating on the decision in front of you right now.

FAMILY 1 The shortcuts that move your numbers

Anchoring

The first number you see pulls the estimates that come after it, including numbers it has nothing to do with.

SHOWS UP

When a proposal arrives with a previous figure already attached.

Framing

The same option worded as a gain and as a loss gets different answers from the same population.

SHOWS UP

When a case states the upside of acting and leaves the cost of acting implicit.

Overprecision

Your confidence ranges are narrower than your accuracy justifies, and the gap does not show until you score them.

SHOWS UP

Any time you give a single number, or a range you have never scored.

Planning fallacy

A schedule built from this project's own details describes the best case, not the likely one.

SHOWS UP

Every plan built from the bottom up, by the people who will do the work.

FAMILY 2 The pull to be right

Confirmation

We go looking for the evidence that agrees with us, and agreement feels the same as accuracy.

SHOWS UP

Whenever you already have an answer you would prefer.

Self-serving attribution

Wins get explained by our judgment and losses by circumstances, so neither teaches much.

SHOWS UP

In any review that happens after the result is known.

Escalation of commitment

What you have already spent raises what you are willing to spend next, which is backwards.

SHOWS UP

At the second and third funding decision on the same programme.

Bias blind spot

You see all of this in other people first, and status raises confidence without raising accuracy.

SHOWS UP

Whenever you are sure you are the objective one in the room.

FAMILY 3 What the room does to a judgment

Going along

About a third of answers bend to a confident but wrong majority, and almost nothing predicts who bends.

SHOWS UP

In any room where other views are stated before yours.

Who speaks first

The first view stated becomes the one everyone after has to argue against, and rank makes it heavier.

SHOWS UP

Sign-off chains, panels, and any meeting a senior person opens.

Shared information

Talk dwells on what everyone already knows; the fact one person holds is the one most likely never said.

SHOWS UP

When the decision turns on something only one person saw.

Silent dissent

A leader who signals a preference early, and a minority that decides it is not worth speaking.

SHOWS UP

When a decision meeting is unusually smooth.

What helps: collect judgments privately before the meeting, put a number on your calls and score them, write the reasoning before you know the result, and ask where each inherited figure came from. **From** Catching your own blind spots, a free three hour course at Overtly Learning, where every claim is cited to an open source and every countermeasure says whether anyone has tested it.

Exhibit 1 is also a standalone one page download, for printing and for putting in front of a team.

The output still looks like arithmetic, which is what makes this family hard to see.

Where the estimate starts. Anchoring is leaning too heavily on one piece of information. In the classic demonstration, people spun a wheel of fortune and then estimated the share of African countries in the United Nations: those whose wheel stopped at 10 said 25 percent, those who got 65 said 45. A number carrying no information moved the answer.

At work the anchor often carries information, which makes it harder to leave. Hong Kong's MTR Corporation had a strong record building urban and conventional rail. Asked to build the territory's first high-speed line, it planned from that record rather than from what high-speed rail costs elsewhere. The competence was real. The starting point was the problem.

Two things make this one uncomfortable. You cannot screen for it: anchoring tasks measure individuals so unreliably that picking out who is susceptible is not supported. And experience does not release it. Buyers whose job is evaluating suppliers were shown a supplier's past cost score, and it pulled their ratings on quality, service and delivery, none of which the score was about.

How the option is worded. Put an epidemic policy choice to World Bank staff. One option saves a third of the population for certain; the other is a gamble with the same expected result. Described as lives saved, 75 percent took the certain option. Described as deaths, 34 percent did. Same numbers, same population, different answer. These were people who set policy for a living.

There is a practical handle in the detail. The effect lives in how completely the uncertain option is described: delete one half of the gamble and it largely disappears, delete the other and it grows. Most proposals at work arrive described one way round, with the other half left implicit. The half that goes unstated is the half doing the work.

A range that is never wrong is too wide. A range that is often wrong is telling you something. A range nobody has checked is telling you nothing at all.

How wide the range is. Asked for ranges wide enough to catch the right answer nine times in ten, people produce ranges that catch it between two and four times in ten. It is not a fixed human constant: across four populations it depended on the task, the incentives, the feedback and the culture, and was sometimes weak and sometimes reversed. Expect yours to be too narrow, then find out.

Reviewing the Challenger disaster, Richard Feynman found that NASA could not agree on how likely a Shuttle was to fail. The working engineers put it near 1 in 100. Management put it at 1 in 100,000, which means launching every day for 300 years and losing one. Two numbers, three orders of magnitude apart, inside one organization.

How the schedule is built. A plan built from the details of this project is the inside view, and it describes the best case. Across eight categories of capital project, average costs exceeded estimates in every single category, from 1.24 times for roads to 1.96 times for dams.

A team in Israel writing a new school curriculum each estimated how long it would take, and the answers ran from 18 to 30 months. Someone then asked the curriculum expert in the room how long similar teams had taken: about 40 percent never finished, and he could recall none that finished in under seven years. The team carried on. It finished eight years later. The outside view was in the room, it was said out loud, and it changed nothing.



3 · THE PULL TO BE RIGHT

Agreement feels like accuracy, which is the whole problem.

Seeking what agrees. We look for the evidence that supports the view we already hold, and the sensation it produces is indistinguishable from being correct. The countermeasures that asked people to reason harder have no clean result behind them. The one that worked did something else entirely: it showed people which advisors had actually been right, and they chose accordingly.

Claiming the wins. When a call goes well we explain it with our own judgment; when it goes badly we reach for circumstances. The World Bank's evaluation group notes that its own task team leaders claim more responsibility for successes than for failures. Being made to justify a judgment afterwards does not correct this, which is why the check has to happen before anyone starts assigning causes.

Spending because you have spent. Money already gone should carry no weight in what you spend next. In practice the opposite holds, and the more a programme has absorbed the harder it is to stop without looking wasteful. Economists at a retreat answered a forestry case alone and then again in pairs: alone, a higher sunk cost raised their likelihood of committing more funds by 15 percent, and in pairs the effect was gone. That is one small study, which is worth knowing before anyone redesigns a governance process around it.

Seeing it in everyone else. Status and power raise confidence without raising accuracy, and rank changes what people believe about their own organization. In one staff survey, 41 percent of senior grades said mistakes were learned from, against 17 percent further down. Same organization, two different pictures of it.

Rank tells you who is trusted. Only a record of what turned out right tells you who to believe.

EXHIBIT 2

How far the numbers move, in three measured findings

None of these is a story about carelessness. Two were measured on tasks with a known right answer. The third is the record of 2,062 finished projects.

25% / 45%

The same question, answered after a wheel of fortune stopped at 10 and at 65. A number carrying no information moved the answer by twenty points.

wheel at 10, answered 25%

wheel at 65, answered 45%

Estimating the share of African countries in the United Nations.

2 to 4 in 10

People were asked for ranges wide enough to be right nine times in ten. This is how often those ranges caught the answer.

intended 9 in 10

achieved 2 to 4 in 10

It is not a human constant: across four populations it depended on the task, the incentives, the feedback and the culture.

8 of 8

Categories of capital project where average cost came in above the estimate, from 1.24 times for roads to 1.96 times for dams.

roads 1.24x

1.0x on budget

dams 1.96x

scale 0 to 2.0x

The pattern runs one way in every category measured.

From Catching your own blind spots, Overtly Learning. Sources are cited in the course and in the paper this exhibit belongs to.

4 · WHAT THE ROOM DOES

A judgment can be sound when it walks into a meeting and not when it walks out.

Going along. Put a person in front of a confident majority giving an obviously wrong answer and about a third of answers bend. Almost nothing predicts who bends: of the personality traits measured, only one showed a clear link. Those studies used strangers rather than colleagues, which is not a reason to assume you are exempt.

Who speaks first. The first view stated becomes the thing everyone after has to argue against, and rank makes it harder still. On the eve of the Challenger launch, Thiokol's engineers found the burden of proof reversed: it was now on them to prove it was not safe to fly. The order of speaking is a decision in itself, and usually one nobody made deliberately.

What everyone already knows. Group talk gravitates to shared information, because more people hold it and because we rate people as more competent when they know what we know. The fact only one person holds is the one most likely never to be said. Data showing that Hubble's primary mirror was flawed existed while the mirror was being made; it was not recognized or fully investigated, in a contractor division working in what the investigation called a closed-door environment.

What nobody says. The famous groupthink story blames cohesion, and when that was tested nothing supported it. What does damage a decision is a leader who signals a preference early and a minority that decides it is not worth speaking. Before the Mars Climate Orbiter was lost, working-level staff had concerns about gaps between navigation solutions; when data conflicts surfaced the team used email rather than the formal problem-reporting process, and the problem slipped through the cracks.

Change the process, not the person.

Training people out of bias is the obvious move and has very little showing it reaches real decisions. Only about a dozen studies have checked whether such training survives even two weeks. Almost all used games or videos, most of them still worked after the gap, and exactly one tested whether the effect transferred to anything else.

The changes that act on how a decision is made have better records. The clearest result: clinicians who estimated alone and then saw the average of four unnamed peers improved twice as much as those who revised alone, and the least accurate improved most. That was on case vignettes rather than real patients, in a network where nobody had more sway than anyone else. Both are conditions rather than details.

Two more are worth knowing precisely because of what is not behind them. Putting somebody in the role of challenger is widely recommended and has barely been tested as organizations actually practise it; the one formal version anybody has tested, a matrix for weighing competing explanations, came back with little or nothing. The pre-mortem, where a team assumes the plan has failed and works backward to why, is recommended everywhere and nobody has compared a plan that had one against a plan that did not.

Four things to do this quarter

1. **Collect judgments before the meeting, not in it.** Each person writes their estimate privately; everyone sees the unnamed average; then you discuss.
2. **Put a number on your calls and score it.** Five decisions that will resolve within three months, each with a probability written down beforehand.
3. **Write the reasoning before you know the result.** The decision, the options, two sentences on why, and a review date. Judge the reasoning on that date before you look at the outcome.
4. **Ask where the number came from.** For any figure you inherit: where did this come from, and was that thing like this thing?

None of these makes anyone unbiased. The most a technique does is make a particular mistake less likely, in a setting somebody has checked. That is a smaller promise than the training industry makes, and it is the one the evidence will carry.

The full course, **Catching your own blind spots**, is free and takes about three hours. Every claim in it is cited to an open source, and every countermeasure says plainly whether anyone has tested it.

What bias does to a decision. Published by Overtly Learning, October 2026. Written by Orlando Franco, strategic advisor.

Every claim in this paper is carried by the corpus of the free course **Catching your own blind spots**, where each is cited to an open, published source. Exhibits 1 and 2 are built from that corpus. The course takes about three hours and costs nothing.

Use and share this paper as it is. Not for resale. The cited works stay under their own licences.