SUBSTACK 2026

The advice was good. It didn't matter.

Dr Leon Rozen

There is a particular kind of professional frustration that doesn't have a clean name. It isn't failure — the work was sound. It isn't incompetence. The reasoning was careful, the constraints were weighed, the recommendation was defensible. It is something more specific: the experience of having done good work and watched it make no difference.

The recommendation was heard. It wasn't listened to. Or the response provided the appearance of consultation without the substance of influence. Or the decision went another way entirely, and nobody could quite explain why in relation to the recommendation given.

This happens to capable professionals with some regularity. It happens in organisations that would describe themselves as evidence-based, data-driven, cross-functionally integrated. It happens to people who are, by any reasonable assessment, doing their jobs well.

I've been thinking about why for a long time.

The standard explanations get offered quickly and with some confidence. Politics is the most common. The decision was already made, the process was theatre. Sometimes this is true. But it doesn't explain why the same pattern recurs across different forums, different decision-makers, different organisations. If it were just politics, you'd expect more variation.

Presentation is another favourite. The argument wasn't framed compellingly enough, the ask wasn't clear, the slides weren't structured right. There's an entire industry built on this explanation — executive presence coaching, communication skills training, influence frameworks that promise to close the gap between good thinking and organisational impact. Executive presence coaching alone is a multi-billion dollar industry.

The problem is that capable professionals who address the presentation problem find the underlying pattern persists. Better slides, clearer asks, more compelling delivery does not always fix the gap. At some point it becomes worth asking whether the diagnosis is right.

This morning I sat in a room where the discussion was animated, the stakes were real, and the decision was genuinely unresolved. I had seen what I thought was a clear path through. I took some care with how to frame it. The words, the sequencing, the timing. I said my piece with some confidence and waited.

The room went silent for a moment. Then moved on.

I have no idea whether what I said was brilliant or irrelevant. Either response would have looked the same. That ambiguity — and the complete opacity of the evaluation — is not a presentation problem. I had done what the standard advice recommends. The silence was confirmation that, whether brilliant or irrelevant, the weight my contribution would carry had been allocated by something I could neither see nor influence.

This is not an isolated experience. And I suspect it isn't for you either.

My view is that the diagnosis the standard explanations offer is wrong. The gap between reasoning quality and decision influence isn't primarily a communication problem. It's a structural one, produced by a specific mechanism that operates in senior decision environments regardless of the quality of the individuals involved or how well they

present. Understanding that mechanism doesn't make the gap disappear, but it changes what kind of problem it is. And that changes how we need to approach it.

The mechanism begins with a constraint that no organisation has solved: senior decision-makers cannot directly observe the quality of the reasoning behind an advisory contribution. They can observe the output — the recommendation, the framing, the position taken. They cannot observe the process that produced it. They cannot see which constraints were held simultaneously, how the trade-offs were weighed, where the uncertainty was genuine and where it was residual ambiguity that further analysis could resolve.

This is not a failure of communication. Senior advisory roles require holding competing constraints across scientific, commercial, regulatory and operational domains and producing a defensible directional stance from them. That kind of reasoning is inherently tacit. It doesn't transfer intact into any deliverable, however well constructed. That's a property of what expert judgment actually is, not a problem that better communication solves.

In a world without time pressure, decision-makers could probe deeply enough to partially reconstruct the reasoning behind any contribution. Senior decision forums are not that world. Decision-makers are working across multiple sources and domains simultaneously, under time pressure, with incomplete information and finite cognitive bandwidth. The human brain, under those conditions, can't evaluate everything — so it filters. It looks for signals that match established patterns of what reliable reasoning looks like, and it discounts what doesn't fit. That filtering is efficient, and usually good enough. But it is systematically distorting in ways that follow predictable patterns, and those distortions fall harder on some advisory functions than others.

The gap, in other words, is not random. Professionals who carry this frustration are not uniformly distributed. They cluster in functions that face a particular combination of structural features: the knowledge in their domain is deep and complex; their value on

a day-to-day basis is counterfactual; and the institutional history of the function has biased the lens through which their contributions are received.

When these features converge, the filtering mechanism is working from distorted inputs. The gap between reasoning quality and what gets attributed in terms of decision influence widens. Not because the reasoning is flawed, but because the mechanism that converts reasoning into influence is operating on inadequate information.

That is a diagnostic claim. It says the gap has a specific cause, the cause produces predictable patterns, and those patterns can be identified. Which is the necessary precondition for any intervention that might actually work.

A few things worth saying at the outset about what this channel is not.

It is not about executive presence. Executive presence is a construct about performing authority — signalling the surface features that audiences associate with leadership. That is precisely the wrong frame. The problem isn't that advisory professionals aren't projecting the right signals. It's that the filtering mechanism decision-makers rely on is operating on proxies that may bear little relationship to the actual quality of the reasoning being offered. Those are different problems. They have different answers.

It is not self-help. The mechanism operates at the level of decision systems, institutional history, and the structural features of expert cognition. Individuals are caught in it. They didn't produce it and can't opt out of it by attitude adjustment.

And it is not a framework for navigating organisational politics. I am not particularly interested in helping people become better political operators. I am interested in the more specific question of why reasoning quality and decision influence come apart — and what that implies for professionals who want to close that gap through legitimate means.

The ideas developed here come from a body of work I've been building for quite a while, grounded in established research and applied to the problem of decision influence in

senior advisory contexts. The channel will develop that work one construct, one mechanism, one argument at a time.

We start with the gap itself because without a clear account of what is actually happening, everything else is just noise.

Dr Leon Rozen · leon@medicalaffairsmentor.com