



RSB Bank

Right Answer, Wrong Audience



Background

Case

RSB Bank — Rangoon Singapore Banking Corporation — traces its origins to the 1860s, when it was founded to finance the growing trade between British merchants and the ports of Southeast Asia. What began as a handful of trading-houses in Rangoon and Singapore grew into one of the world's largest banking and financial services groups. Today, headquartered in London and led by CEO Nigel Ashworth, RSB's commercial, investment, retail, and private banking divisions operate in more than 80 countries. The Financial Stability Board (FSB), the international body established by the G20 to monitor risk in the global financial system, designates RSB a systemically important financial institution — one of 29 banks worldwide considered too big to fail.

The 2008 Global Financial Crisis put banks like RSB under a level of regulatory scrutiny they had never faced before. In the years that followed, industry regulators trying to understand what had gone wrong kept arriving at a similar diagnosis: money laundering, misleading investors, misrepresenting the risks of financial products, and outright market rigging - all traced back to insufficient risk management. Regulators identified this as what they called a 'weak risk culture': a work environment marked by a lack of risk awareness, in which employees were (implicitly or explicitly) incentivized to take on too much risk. For this

reason, the FSB and other regulators began advising banks in 2011 to build a 'sound risk culture', not merely stronger risk controls on paper. When the UK's Financial Conduct Authority (FCA) formally raised risk culture as a supervisory priority in 2013, RSB's board understood the signal for what it was: regulators were no longer treating this as optional.

RSB's response is to launch a company-wide Risk Culture Program. Elaine Whitcombe, the board member responsible for HR, takes overall responsibility for it, and brings in Meera Anand, an external consultant who previously led similar work at another major bank, Barclays, to run the program day to day. Meera is paired with Tom Whitfield, RSB's own Head of Organizational Change and Development, and the two lead the program jointly, both reporting directly to Elaine. Together, they have to decide where and how to start — they settle on Market Operations.

At RSB, every trade executed by investment bankers in the 'front office' — the client-facing trading desks — must pass through Market Operations, which confirms the trade, logs it into the bank's systems, and moves the associated funds across accounts. Senior executives in RSB's Investment Banking division had grown increasingly concerned about the number of transactional errors within Market Operations. The errors took various forms: money posted to the wrong account, a client charged more than the agreed fee, a trade left unconfirmed or never booked into the system at all. One error type stood out as particularly costly — a trade recorded with the wrong settlement date, which, if the exchange rate or share price moved in the intervening period, could turn a minor slip into a real financial loss.

Whenever an error was serious enough, an incident report was written up. Market Operations executives, reading back through these reports, noticed two things that troubled them. First, most of the incidents looked preventable in hindsight. Second, the reports rarely showed any evidence that lessons were learned — no follow-up, no attempt to adjust the process so the same mistake wouldn't recur. A few of the executives raising these concerns had previously worked at other major banks — Credit Suisse, Morgan Stanley. At those banks, they claimed, an error like this would have triggered an alert, a review, and a concrete follow-up that fed back into how the process worked — a healthy culture of learning and improvement, not just error correction. At Market Operations, however, such a culture was notably absent. Which made this division, in their view, exactly the right place for RSB's new Risk Culture Program to begin.

Step 1. Defining the problem

Question

Before deciding on any solution — in this case rolling out a Risk Culture Program — the problem itself needs to be clearly defined. Assess whether that is currently the case. In particular, consider:

- 1. What is the problem to be solved?*
- 2. How clearly defined is the problem (what, who, when, where, why)?*
- 3. What assumptions underlie how RSB's executives have framed the problem?*
- 4. Is it clear what the organizational consequences of the problem are?*
- 5. Is it clear how serious and urgent the problem is?*
- 6. Is it clear what the logic model (causal mechanism) is: what are the underlying cause(s), how do they lead to the problem, and how does the problem consequently lead to the organizational consequences?*

Your answer

Step 2. Identifying evidence gaps

Case

Meera and Tom sit down for their first real working session — just the two of them, before anything goes to Elaine. Their assignment is simple on paper: roll out the Risk Culture Program in Market Operations. But Tom doesn't look convinced. "I keep coming back to the same thing," he says. "We don't actually know what the problem is. We know what — errors, recurring, some of them costly. But who or what is actually responsible, whether it's getting worse or better, why it's happening — none of that's been established. It's mostly a handful of executives comparing what they've seen at RSB with their experience at banks they used to work at. In fact, the whole idea that it's a risk culture problem — that's an assumption. It could just as easily be understaffing, unclear procedures, or the systems themselves."

Meera listens, turning her pen over. "It's true no one's really tested that assumption — some of these executives may just be repeating what they've heard from others," she says. "But the board seems pretty convinced. And an experienced group of executives converging on the same explanation isn't strong evidence on its own, but it's not something to dismiss either." Tom shrugs. "I agree it's a reasonable working hypothesis. But I think it's on us to close the gaps in what we actually know first. If we can't even say clearly what the problem is, I don't see how any program — this one or any other — could confidently claim to fix it."

Question

The case so far showed several uncertainties regarding the assumed problem. As Tom says, these uncertainties need to be resolved before any solution can be considered. We therefore need to determine:

- 1. What evidence, from which sources, is needed to establish whether — and if so, how, for whom, and when — a weak risk culture is contributing to preventable errors in Market Operations?*
- 2. What evidence is needed to establish whether this problem is serious and urgent enough to require action now?*

For each of the four evidence sources (practitioners, the organization, the scientific literature, and stakeholders), describe what you would like to know.


Your answer

Step 3. Acquiring the evidence

Question

In the previous step, you identified what evidence is needed from each of the four sources: practitioners, the organization, the scientific literature, and stakeholders. For each source, describe how you would acquire the evidence (i.e. which method you would use).

Your answer



Step 4. Appraising the trustworthiness of the evidence

Case

Tom walks Meera through what he thinks should be done next: practitioner input via a Delphi panel or a survey, a thorough analysis of organizational data and KPIs pulled from

RSB's own systems, a systematic search across multiple research databases — both broadly, on predictors of human error in procedural work, and more narrowly, on risk culture specifically — and focus groups and a wide survey of Market Operations staff.

Meera doesn't look convinced. "Tom, this is a lot. Surveys, focus groups, literature searches — this could add weeks, maybe months, before we've done anything. The board wants this program rolled out. That's the assignment." "I understand," Tom says. "But if we roll this out without knowing whether we're even solving the right problem—" "I know," Meera cuts in. "But that's not my call to make. I was brought in to do essentially what I did at Barclays — take a Risk Culture Program and roll it out. Nobody asked me to re-open the question of whether the diagnosis is correct. If you think that needs to happen, that's a conversation for you to have with Elaine, or the board. I won't be the one pushing that upstairs. I can't renegotiate my own assignment."

Tom sits with that for a moment. He knows she's right, but he knows this isn't a battle he's going to win. "Okay," he says slowly. "I agree we can't do all of it. But some of this is easy — the organizational data, we can pull that ourselves, it already exists. A short survey to the most experienced managers, open questions, that's not a heavy lift either. Regarding the literature review — I guess a few well-chosen searches in Consensus could suffice, that's a matter of days, not weeks. Stakeholders are the harder sell. But that's also the evidence I think matters most. They're the ones actually doing the work, so they have the most relevant insights. So maybe no survey or formal focus group — maybe just spend some time on the work floor, have casual conversations, see what comes up. Lighter touch, but still something."

Meera considers it. It's not what she'd have chosen, but it's a compromise she can live with. They settle on a split: Tom and his team take on the evidence-gathering, while Meera turns her attention to preparing the program rollout itself.

Question

It is clear that Tom has made some serious compromises regarding what (content) and how (method) evidence will be acquired. What do you think of his choice to conduct a short survey with open questions to the most experienced managers, rather than the Delphi panel originally proposed? How trustworthy would you consider the outcome?

Your answer

Question

How would you judge the trustworthiness of findings from meta-analyses and longitudinal studies identified through a search with a tool like Consensus AI?

Your answer

Question

How would you judge the trustworthiness of stakeholder evidence acquired by spending some time on the work floor and having casual conversations with the staff within Market Operations?

Your answer

Step 5. Aggregating the evidence

Question

The first evidence Tom and his team obtain concerns the findings from the short survey with four open-ended question from 12 experienced operations, risk, and compliance managers. Download a report with the results << here >>. Summarize the 12 responses and make a list with the main findings. You can use an AI tool such as Claude or ChatGPT to support you.

Your answer

Case

When Tom goes through the organizational data he received, he notices that it is much thinner than expected. The error and trend data exist, but can't be linked to staffing, workload, or any HR indicator. This is because mistakes were never logged against individual employee numbers, so there's no way to connect who made an error — on which team, at which location — to any metric in the HR systems. The cost data tell a similar story: only errors above the £20k threshold were logged, and even then, only for client compensation or exchange-rate losses tied to date errors — there is no data available on fees or penalties, in fact, there is no information on errors below the threshold at all. In addition, the incident reports don't appear to be centrally held; as far as Tom can tell, they sit wherever individual managers happened to keep them. Finally, the one comparison Tom had been counting on — benchmarking RSB's error rate against the rest of the industry — turns out to be a dead end too: this kind of data is treated as competitively sensitive and isn't shared across banks. What looked, on paper, like the most straightforward evidence source to acquire has turned out to be the most disappointing.

Question

Download the error rates & costs dataset <<here>>.

- A. What are the total annual costs associated with the errors?*
- B. Is the error rate, and associated costs, rising, falling, or stable over the four-year period?*
- C. What does this tell you about the seriousness and urgency of the problem?*

You can use an AI tool such as Claude or ChatGPT to support you with the calculations.

Your answer

Case

When Tom has a second look at the error rate & costs file, something catches his eye. Skimming down the monthly error counts and associated costs, broken out by location, the numbers aren't evenly spread — some locations sit noticeably higher than the rest, month after month, while others stay consistently low. This is not a one-off spike either; the pattern holds steady across the whole period covered in the file.

Question

Calculate the average error rate and average associated cost separately for the 'Onshore' and 'Offshore' groups. Then calculate whether the offshore group's average error rate and average associated cost are higher compared to the onshore group, expressed as a percentage. You can use an AI tool such as Claude or ChatGPT to support you with the calculations.

Your answer

Case

Tom stares at these numbers for a while. Six locations, consistently three and a half times higher than the other three, for years running. It's the kind of pattern that's too clean to be random, but he can't yet say what it actually means. Then he remembers that one of the managers in the survey made a comment about offshore locations. He makes a mental note to discuss this next meeting with his team.

He then turns to the next piece of evidence - the summary of the review of empirical studies on the impact of (risk) culture and the summary of the review of empirical studies on predictors of human errors. When he reads through the main findings, he is visibly pleased. "I think we've actually got something here."

Question

Download a copy of the summary on the impact of risk culture <<here>> and a copy of the summary on predictors of human errors <<here>>. Read both summaries - what do you consider the three most important insights for the case?

Your answer

Case

The last piece of evidence concerns the stakeholders. Tom spent a full day on the floor at Market Operations, RSB's London headquarters, moving between desks and taking the opportunity to talk informally with most of the staff. At that point, the findings from the practitioner survey, the organizational data, and the scientific literature were not yet available to him. He considers this both an advantage and a disadvantage: it let him

approach the conversations with an open mind, unanchored by what the other sources might suggest — but it also meant he could not probe deliberately into any specific patterns those sources revealed.

Question

Download a summary of Tom's interviews with the staff <<here>>. Compare the report against the Delphi panel findings, the organizational data, and the scientific evidence reviewed earlier. On which specific points does the evidence converge?

Your answer

Question

Now that you've cross-validated the four evidence sources, what does this convergence tell you about the strength of the aggregated evidence?

Your answer

Case

With the evidence gathered, Tom considers bringing representatives of each source together — practitioners, stakeholders, and the people who reviewed the scientific literature and organizational data — into a cross-validation session. But time is against him, and pulling people from London, Boston, and six offshore sites into one room or a virtual session isn't realistic on the timeline the board expects.

Instead, he brings the findings to his own team: twelve experienced change managers, all trained in evidence-based decision-making. He walks them through each individual source — the practitioner survey, the organizational data, the scientific literature, the staff interviews — and then presents the seven aggregated points. In the discussion that follows, some members ask for clarification, a couple of minor comments are made, but no real pushback. The team arrives at the same conclusions — the seven points hold up.

Then Tom brings up something that had been sitting at the edge of the discussion: the earlier comment from one of the survey respondents about the offshore locations. A couple of the managers around the table start to chuckle before he's even finished the sentence. One asks, not unkindly, whether he's ever actually been to one of those sites. He hasn't.

The team explains that the six offshore locations bear almost no resemblance to London, Boston, or Singapore. Most are large, packed halls — hundreds of people working shoulder to shoulder, in loud, chaotic, physically rough conditions. Genuinely, as one manager puts it, brain-hostile environments. Staffing is its own separate crisis: recruitment in these locations is so under-resourced that people are hired with barely any screening, minimal qualifications, no real onboarding, and only a fraction of the training London staff receive.

Put next to the seven points the group had just agreed on — in particular interruptions & distractions, workload & time pressure, and inexperienced & undertrained staff — the picture sharpens considerably.



The offshore sites share the same problems as London, but in a dramatically worse form, on every dimension the evidence has identified. For the first time, Tom has an actual explanation, not just a correlation, for why the offshore locations show error rates and costs three to four times higher than the onshore locations.

Step 6. Applying the evidence

Question

Distractions and task interruptions appear to be an important contributor to errors within Market Operations. What would you recommend to the board?

Your answer

Case

With the interruption findings and possible solutions sketched out, Tom moves the discussion with his team to the remaining points: the workload spikes, the outdated confirmation system, and the inexperienced, undertrained staff. On workload spikes, there's fairly quick agreement: smoothing out the sudden clustering of confirmations, perhaps through different scheduling or short-term surge staffing during peak periods, seems achievable without major structural change. The confirmation system itself is a bigger undertaking — nobody at the table thinks replacing the IT infrastructure is a quick fix, but it's a well-understood kind of project: budget, timeline, vendor selection, the usual. Similarly,

for RSB, training and onboarding is a 'routine' issue — it's simply a question of investment and priority, not of figuring out something new.

Applied to London, Boston, or Singapore, the team is reasonably confident that these interventions are feasible. But when the conversation turns to the six offshore locations, the confidence drops. The scale of the gap described earlier — the packed halls, the barely-screened hiring, the near-total absence of onboarding — isn't something a few interventions can close. It's at this point that one of the more senior managers says what several others are already thinking. "We should be careful how we frame this to the board," she says. "This isn't just an operational finding. It's a direct critique of RSB's offshoring strategy — and that strategy has Nigel's name on it." Everyone in the room knows what she is referring to. CEO Nigel Ashworth was brought in specifically to cut costs, and moving functions like confirmation processing to lower-cost countries has been one of the most visible expressions of that strategy since he took over. Concluding that this same strategy is now a major driver of RSB's error and cost problem isn't a neutral observation — it's a direct challenge to a decision Nigel has publicly staked his tenure on.

Step 7. Assessing the outcome

Case

The team is now left with one final question: if the board were to accept the proposed interventions, how could it be assessed whether this reduces the error rates and affiliated costs? The error rates and costs are measured monthly — so there is a baseline and changes can be reliably monitored. The best method to assess the impact would be to randomize four locations to a group where the interventions will be implemented, and compare them with the other five locations that continue their operations as usual. Clearly, this would not be a feasible approach.

Question

If a randomized controlled trial is not possible, what would be the best alternative method to assess the impact of the interventions?

Your answer

Case

Two weeks later, Tom is invited to sit in on a board meeting where Meera is scheduled to present the design of the Risk Culture Program. She walks the board through RSB's new Risk Culture Framework: five core values — risk awareness, risk behaviour, risk ownership, risk leadership, and supportive environment — each broken into four or five underlying elements, each of those translated into concrete behavioral expectations. For example, under risk awareness, the element "role-specific risk awareness" is operationalized as *"All staff understand and follow their procedures. They demonstrate an understanding of the risks associated with their role and their impact."* Implementation, she explains, will run through nine working groups — one per location — involving roughly 50 people worldwide, who will lead six weeks of group discussions with local staff to embed the framework.

The board has a handful of minor questions — timeline, cost, the usual stuff — but overall they seem content. Several board members compliment the clarity of the presentation. Meera thanks the board and leaves the room.

Once the door closes, Elaine Whitcombe speaks. "Before we move on," she says, "I've had some serious conversations with Tom over the past few weeks, about some findings I think this board needs to see."

Tom takes the floor and starts his presentation, walking the board through the four evidence sources: the practitioner survey, the organizational data, the scientific literature, the staff interviews. He shows where they converge. He ends on the aggregated findings, starting with the one that matters most for what the board had originally asked for: none of the evidence points to a weak risk culture as a driver of the errors. The likelihood that Market Operations' errors are being caused by a weak risk culture, he tells them plainly, is extremely small.

What the evidence does point to, consistently and from multiple independent sources, is something else entirely: task interruptions and distractions, workload spikes concentrated around clustered trading activity, an outdated confirmation system, in combination with understaffing and inadequate training. When Tom finishes, the room is quiet for a moment.

Then CEO Nigel Ashworth speaks — and what he says catches Tom off guard. He tells Tom this is excellent work, thorough, exactly the evidence-based rigor he expects from his

top managers. Elaine, he admits, had already briefed him in outline, so his findings aren't a total surprise. And, he adds, the board already noticed that the organizational data showed that something was wrong with the offshore sites. Local managers had been sending word up through informal channels for months, about working conditions and recruitment practices that fell well short of what London or Boston operate under. So, he ends, the board had already half-suspected this was behind the elevated error rates at those locations. But thanks to Tom, the board now has the evidence to substantiate their suspicion.

Tom hadn't expected any of this — not the praise, not the acknowledgment that the board had already been circling the same conclusion. It encourages him to ask the question that's been sitting at the back of his mind since Meera left the room. "Can I ask — given everything we've just gone through, what does this mean for the Risk Culture Program?"

There's a beat of silence, and then, unexpectedly, laughter — quiet, knowing chuckles moving around the table. Nigel breaks into a broad smile. "Fair question," he says. Let me ask you one back, though. Who do you think are the most important stakeholders this board answers to — other than our shareholders?" Tom hesitates, then offers the obvious answer. "The regulators." "Exactly. The Financial Conduct Authority. The Bank of England. The Financial Stability Board. And bodies like these have made it unmistakably clear, for over a decade now, that a sound risk culture is treated as central to how banks like ours are governed and supervised — whether or not the underlying science agrees with them." He lets that land for a second. "And you've just shown us, quite convincingly, that it doesn't. No agreed upon definition. No reliable way to measure it. Not evidence-based." He shrugs. "You're not wrong, Tom. But you'll understand, I hope, that this board is not in the business of educating our regulators on the limits of their own frameworks. Being right, and being seen as clever about being right, in front of the people who license us to operate — that would make our shareholders very unhappy.

So what do we do instead? Exactly what we just decided. We run the program. We hold the workshops. We produce the posters, the nice presentations. And when the regulators come to audit us, we show them exactly what they're expecting to see, and we tick the box. Because it satisfies the people whose approval keeps this bank in business." He pauses. "What you've given us — the evidence-based findings — that's what gets acted on, silently, under the radar. But the risk culture program is what gets shown."