

Rooms AND Reasons

ON THE EVIDENCE





Why this book exists

Most of what facilitators know arrives as craft. Do this, and the room opens up. Try that, and the quiet person finally speaks. Sit the group in a circle rather than in rows. Hand out the pens before you ask the question. This knowledge is passed from one practitioner to the next, it is refined by experience, and it usually works.

Almost none of it comes with a reason attached.

The craft without the reasons

That is a strange situation for a profession to be in. We make decisions all day about how people will think together, and we make most of those decisions from habit. When something works, we repeat it. When something fails, we usually blame the room. We rarely ask what is actually happening inside a group of people when they agree too quickly, or go quiet, or become strangely polite, or stop trying.

Researchers have been asking exactly those questions for more than sixty years. They have run experiments, followed organisations for decades, watched hospitals and schools and factories, and published careful findings about how groups behave and why. A great deal of it

is directly relevant to what happens in a room you are running on Tuesday morning.

Very little of it reaches us.

Why the research
does not reach us

There are two reasons for that. The first is that research is written for other researchers. It is dense, it is cautious, and it is full of language designed to survive peer review rather than to be understood quickly. The second reason is more uncomfortable. Many practical people have decided that theory is a waste of time. The word itself makes them close the page. They have met too much theory that arrived as jargon, offered no evidence, and asked to be believed.

That reaction is fair, and it is expensive. It leaves us working from intuition in a field where intuition is often precisely wrong. The groupthink chapter in this book turns on a finding that should worry every one of us. Groups that made disastrous decisions reported high satisfaction with their own process. The feeling of a good session and the fact of a good session came apart completely, and not one person inside the room could tell the difference. If your only instrument is your own sense of how the room went, that finding tells you your instrument is broken in exactly the situations that matter most. There is a second awkwardness for anybody who has decided that theory is somebody else's business. Our own professional standard has already settled the question.

The IAF Core Competencies set out six areas that define capable facilitation. The fifth is Build and Maintain Professional Knowledge. Its first component asks for knowledge in the theory and practice of group process facilitation, and then lists what counts as related knowledge:

organisational development, psychology, conflict resolution, dynamics of change, learning and thinking theory

IAF Core Competencies, E1

Psychology. Learning theory. The dynamics of change. These are not optional extras that a particularly studious facilitator might take up in the evenings. They sit inside the definition of the job, in a framework built over several years by facilitators from all over the world, and they are assessed when somebody applies to be certified.

Which means the question was never whether to engage with the research. The question is only whether we do it well.

What evidence
gives you that
experience cannot

This is what evidence gives you that experience alone cannot. Experience teaches you what usually happens. Research tells you what is actually causing it, which is the only thing that lets you change it on purpose. Experience tells you that adding a competition raises the energy in a room. Research tells you what the competition is quietly doing to the

quality of what people say, and under which specific conditions, and what you can do instead. One is a pattern. The other is a mechanism. You can only design with a mechanism.

Working closely with the research does something else as well. It tells you where the evidence runs out. Every chapter here names the limits of its own findings, because a facilitator who knows that a study was run on eighteen classrooms in one American school district in 1968 is better equipped than one who has been handed a confident rule. Certainty is easy to sell and it is usually a sign that somebody has stopped reading. The honest position is more useful, and it is this. Here is what we know, here is how well we know it, and here is what you might do with it.

**‘Research that only confirms
what we already do is research
that has been selected rather
than read.’**

You will notice that the five chapters are not especially flattering about facilitators. This is deliberate. Research that only confirms what we already do is research that has been selected rather than read. The findings that are worth the paper are the ones that describe something we do badly without noticing, and

there are more of those than any of us would like.

One thing should be said plainly before you start, because it shapes how the rest of the book is meant to be used.

Nothing here is a prescription. There are no rules in these pages, no model to adopt, no correct way to run a session. Anybody who tells you that a single study settles what you should do on Tuesday has either misunderstood the study or is selling something.

What these chapters offer are provocations. They are material for conversation, for reflection, and for thinking harder about choices you currently make from habit. The protocols at the end of each chapter are written to be specific enough to try, and they are starting points rather than instructions. Change them. Argue with them. Run one badly and work out why it failed in your context when it held in a laboratory.

None of this is a prescription

The research is there to give you better questions rather than ready made answers. A facilitator who has genuinely thought about why a room agrees too quickly will design better sessions than one who has memorised a five step process for managing consensus, even if the second one looks more competent on a slide.

So read these as an invitation to think, and treat every protocol as a draft.

What is in this book


On the Flyntrok website we publish a piece every week in this series, called On the Evidence. Each one takes a single body of research and works out what it means for people who lead, convene and facilitate. Some of them look at why organisations copy each other. Some look at how ideas travel through networks. Some look at why intelligent, articulate people defend the very things they say they want to change.

This book is a synthesis of five of them, rewritten and shortened for people who work in rooms.

Five studies. What each one actually found. Where the evidence stops. And at the end of every chapter, something specific you could run in a session next week.

You do not need any background in psychology to read this. Nothing here assumes you have studied research methods, and every term is explained at the moment it appears. The studies are named and dated so that you can go and read them yourself, and I hope some of you will.

The five chapters move from familiar ground to less familiar ground. Start anywhere. Each one is complete on its own.



A closing warning, offered in good faith. Several of these findings are difficult to unsee. Once you know what the four channels of expectation are, you will catch yourself using all of them before the first coffee break, and you will not be able to stop noticing. I am told this is what professional development feels like.

- Kavi

1 When agreement comes too easily



You have run this session. The group worked well. People built on each other's ideas. By the middle of the afternoon a direction had formed, and when you tested it, everybody agreed. Not a single objection. You closed cleanly and finished on time. The client thanked you on the way out and said it was the most productive day they had spent in months.

Then nothing happened. The decision quietly died. Or somebody told you, weeks later, that they had serious doubts on the day and did not feel able to say so.

That room felt like good facilitation. It had every sign of it. The difficulty is that a room reaching agreement well and a room reaching agreement badly look almost exactly the same from the inside, and that is the problem Irving Janis spent a decade working on.

What Janis found

On 17 April 1961, fourteen hundred Cuban exiles landed at the Bay of Pigs. They had been trained and funded by the American Central Intelligence Agency and told that their arrival would start an uprising. Within three days the operation had completely failed. The room that approved the plan contained the director of the CIA, the most senior military officers in the country, and a president **who was** keen to prove his judgement. By any measure of individual ability, this was an exceptional group of people.

Janis was a psychologist at Yale University. He wanted to understand how a room like that produces a decision like that. His answer was published in 1972 in a book called *Victims of Groupthink*. Close groups under pressure to agree, he argued, will quietly trade accuracy for harmony. The members are not weak and they are not unintelligent. They are doing something very ordinary. They are protecting the feeling in the room.

He listed eight signs. A sense that the group cannot fail. Explaining away the warnings. Assuming the group is on the right side. Dismissing outsiders. Pressure on anyone who hesitates. People silencing themselves. Everyone taking the silence as agreement. And someone, without being asked, keeping awkward information away from the group.

Read that list slowly and notice how many of those signs look, from inside the room, like a session going well.

The most uncomfortable part of Janis's work is what the groups said afterwards. They reported being very satisfied with the process. They described the discussion as open. They believed the decision was the best one available. The feeling of a good room and the fact of a good decision had separated completely, and not one person inside the room could see it.

The competencies are clear about our responsibility here, and the wording repays a second reading:

Provide a safe environment for conflict to surface

IAF Core Competencies, C3

Surface is the operative word. It does not say resolve conflict, or prevent it, or steer the group past it. It assumes there is something down there already, and it makes the conditions for its arrival our job rather than the group's. Most of us read that line as being about managing an argument. It is about producing one.

Where the evidence stops

Janis built his theory from historical cases. He studied three in particular, and they are worth knowing, because each one shows a capable group walking calmly into a result that looks obvious afterwards.

Pearl Harbor was an American naval base in Hawaii. In the months before December 1941, warnings accumulated that Japan was preparing to attack. The commanders there discussed the warnings and settled on the view that an attack so far from Japan was improbable. The attack came on the morning of 7 December and destroyed much of the American Pacific fleet.

The Bay of Pigs is the invasion described at the start of this chapter.

Vietnam refers to the years in which the American government steadily increased its military involvement in a war in south east Asia. Through the middle of the 1960s a small group of senior advisers met regularly and approved one escalation after another, while intelligence reports in front of them indicated that the strategy was not achieving what it was supposed to achieve. The war continued for another decade and the United States eventually withdrew.

Janis read the records of all three after everyone knew how the story ended. That is a genuine weakness, because it is easy to find a pattern in anything once you know the outcome.

When other researchers took the theory into laboratories, the results came back more conditional than Janis had claimed. A review of that work published in 1990 found that the effect depended heavily on two things. A leader who states a preference early. And a group that is cut off from outside opinion. Closeness on its own did less than Janis expected.

That matters for you, because those two conditions are exactly the things a facilitator can change.

Janis revised his own position in 1982 and added practical fixes to the second edition of his book. His instinct was that you cannot ask people to be braver. You have to build the disagreement into the process, so that speaking up costs the speaker nothing.

The protocol:

Last Word

Run this in any session heading towards a decision that matters.

Before the room opens. Agree privately with the senior person, the sponsor, or whoever holds the authority, that they will state their position last. Not never, and not neutrally.

Last. Tell them why. Most will agree, because most have watched a room fold around them and disliked it.

At the start. Name one person, out loud, as critical evaluator for this session. Their job is to argue against whatever the room settles on. Say clearly that this is a role rather than a personal view, and give it to somebody different next time so that it never becomes one person's character. This single move removes the cost of disagreement, because the person disagreeing is following instructions.

At the moment of agreement. Stop. Do not close the item. Say that the room has reached agreement and that you are now going to test it.

The test. Ask one question. What would have to be true for this to be the wrong decision? Take the answers in silence and write every one of them on the wall before anybody is allowed to respond to any of them. The silence is the important part. It stops the group from dealing

with the first objection in a way that discourages the second.

Then break. If your design allows it, hold the confirmation at least an hour later, or the following morning. Janis called this the second chance meeting. Agreement that survives a break is worth far more than agreement reached inside one warm hour.

**Before your
next session**

In your last three sessions, how many ended with a decision that not one person in the room had argued against?



Five chapters is a small sample of a very large body of work.

What sits behind these pages is more than sixty years of careful research into how people behave when you put them in a group and ask them to decide something together. Most of it was never written for us. Almost all of it applies to us. Janis was studying foreign policy. Rosenthal was studying schoolchildren. Menzies was studying a hospital ward. Deci was studying a puzzle. Cohen and Levinthal were studying how firms absorb new technology. Not one of them had a facilitator in mind, and every one of them has something exact to say about what happens in a room you are running next month.

So the most useful thing this book can do is make you want to read more of it.

Read more of it

Read behavioural science. Read the organisational research, the work on groups, the work on decisions and learning and motivation. Read it slowly and read the limitations sections, which are usually the most interesting part of

any paper. You will not agree with all of it. You will find findings that contradict each other and fields that have argued the same question for thirty years without settling it, and that argument is the useful part rather than a sign that the work has failed.

This is not only a professional investment. The research in this book is about how people form expectations of each other, how they protect themselves from difficulty, what happens to their motivation when you reward them, and what a group can and cannot hear. That describes a facilitated session. It also describes a team, a family, a friendship and a marriage. People who understand this material tend to lead better, and they also tend to live with a little more clarity and a little less frustration, which is worth more than any technique.

You will get better at facilitation from this. That is the smaller prize.

2 The ones you have already decided about



Twenty minutes into the day, you have made up your mind about somebody. The one with folded arms. The one looking at their phone. The one who asked a question about scope in a tone you did not like. You have not said anything and you would not describe it as a judgement. You simply know that this person is going to be difficult.

What follows takes the rest of the day, and almost none of it is deliberate.

What Rosenthal found

In 1968 Robert Rosenthal and Lenore Jacobson published a study of a primary school in California. Every child had taken a test at the start of the year. The teachers were told it was the Harvard Test of Inflected Acquisition, and that it could identify children who were about to develop rapidly. They were then given the names of the children in their class who had been identified.

There was no such test. There is no such thing as inflected acquisition. The names on the list had been chosen at random, one child in five, with no connection to anything the children had done. Eight months later every child was tested again, and the children on the list had gained more than the others.

Nothing about those children was different on the day the list was written. The only thing that changed was what their teachers expected of them.

Rosenthal then did the most useful part of the work. He recorded hundreds of hours of teaching to find out how an expectation travels from one person's mind into another person's results. He found four channels, and the teachers were almost completely unaware of using any of them.

They were warmer with the chosen children. More smiles, more eye contact, a gentler voice. They gave those children harder material. They gave them more opportunities to answer, and more time to think before moving on. And they gave them better feedback, with more praise and clearer help when something went wrong.

None of that is a belief. All of it is behaviour. The expectation leaves your head and becomes the conditions the other person is working in.

Every one of those four channels is available to you in a room. You are using all four within the first hour.

There is a line in the competencies that most of us have read as being about other people:

Help individuals identify and review underlying assumptions

IAF Core Competencies, C3

It does not say which individuals. On the evidence in this chapter, the individual in the room whose underlying assumptions are shaping the most outcomes is the one who wrote the agenda.

Where the evidence stops

The original study was criticised heavily and some of that criticism was justified, particularly about the intelligence testing.

The most useful correction came from Stephen Raudenbush in 1984. He combined eighteen studies and found that the effect held strongly under one specific condition. It worked where the teacher had not yet formed their own view of the child. Where a teacher already knew a pupil well, telling them something new about that pupil made very little difference. The effect works on judgements that are still forming, and it fades on judgements that are already settled.

That condition disappeared when the research reached managers, where it arrived as a simple instruction to believe in your people. For you, the condition is the entire point. You meet a room once. Every judgement in it is still forming. You are working in exactly the

window where this effect is strongest.

The research moved into adult settings through Dov Eden at Tel Aviv University, who tested it on Israeli army trainees. A review of seventeen workplace studies published in 2000 found a large average effect. Both found the same shape. It is strongest with new people, and strongest with people who are currently performing badly.

The harder half of this is the same mechanism running backwards. Decide quietly that somebody is limited, and you give them less time to speak, thinner feedback and easier work. They perform accordingly. You turn out to be right, in a way that you caused.

The protocol:

The written off list

This one is done alone, and it is uncomfortable by design.

At the first break of any session, before anything else. Take two minutes. Write down the name of every participant you have already formed a negative view of. Do not soften the entry or explain it. If you have thought it, write it.

Against each name, write what you have actually observed. Not what you concluded. Folded arms is an observation. Disengaged is a story. Asked three questions about scope is an observation. Obstructive is a story. Most entries will not survive this step, and the ones that do not survive are the ones worth looking at hardest.

Then use the four channels deliberately for the rest of the day. For every name on that list, do four things. Give one direct invitation to speak, by name, on something they are likely to know about. Give full attention when they answer, and the same waiting time you give everybody else. Respond to the substance of what they say rather than thanking them for it. And have one conversation at a break where you ask them something and sell them nothing.

You are not being generous here. You are removing a disadvantage that you applied

without deciding to.

After the session. Go back to the list and mark the people you were wrong about. Keep the list across sessions. The pattern in who you write off is far more useful to you than any single entry. It is usually a type of person, and it usually says more about you than about them.

Before your
next session

Who in your next room have you already decided about, based on the briefing call, the delegate list, or their job title?

3 The agenda as armour



The conversation was opening into something real. Somebody had said the thing that everybody knew and had never put into words, and the room had gone very quiet. You could feel where it was heading, and you could not predict what would happen there.

So you moved to breakouts. Or you brought out the framework, or handed round the template, or looked at the clock and said that you needed to keep the group on time.

The structure worked. Structure always works. That is what makes this difficult to see. It was solving a problem that belonged to you rather than to the room, and that distinction is the one Isabel Menzies drew in 1960. It has barely been improved on since.

What Menzies found

Menzies was a researcher at the Tavistock Institute in London. She was called into a teaching hospital to solve a practical problem. There were too many qualified nurses and not enough patients for the student nurses to get proper hands on training. She went looking for a scheduling problem. She found something else.

Student nurses were leaving the programme in large numbers. The training was demanding and the work was meaningful, and they left anyway. What Menzies noticed was how the ward had been organised.

Nurses worked from a fixed list of tasks rather than using their own judgement. Patients were referred to by bed number and condition rather than by name. Tasks were divided so that no nurse spent a long period with any one patient. Responsibility was spread so widely that no single person ever quite held a decision. Checks were repeated and then repeated again, long past the point of any practical use.

Nursing involves an emotional exposure that most jobs do not have. Constant closeness to suffering and death. Decisions that can cost a life. Families arriving frightened and angry. Menzies argued that every one of those practices was a defence against that exposure. The task list meant a nurse could follow instructions instead of being present with somebody's fear. Bed numbers meant there was no person there to lose. Divided tasks meant no attachment could form. Shared responsibility meant that no single nurse carried a

decision alone.

None of this was designed. No meeting agreed it. It had built up slowly, one practice at a time, because each practice made the anxiety easier to bear.

She called it a social defence system. Her sharpest point was this. It worked on the anxiety and did nothing at all about the actual problem. The students who left were not leaving because the work was too hard. They left because the system had made it impossible to do the work well.

The second of the six competency areas is Plan Appropriate Group Processes. It is a reasonable standard and it contains a quiet gap, because appropriate is doing a great deal of work in that sentence and the standard never specifies appropriate for whom. A process that keeps the facilitator comfortable can meet every visible test of good design. It will be well sequenced, properly timed, and thoroughly defensible in the debrief.

Where the evidence stops

This is one researcher, in one hospital, using a psychoanalytic framework, published in 1960. It is observation rather than measurement. The mechanism she proposes is an unconscious institutional defence, which is almost impossible to test directly, because any evidence against it can be read as further defence.

Treat it as a lens rather than as a finding. It has lasted sixty five years because practitioners keep recognising something in it. That is a weaker form of evidence than a replication, and it is a real one.

You may notice that a theory which cannot be disproved is exactly the sort of thing this book is supposed to be sceptical about. That is a

fair objection and I have no good answer to it. I have kept the chapter because the idea has changed how I design sessions more than any properly controlled finding I have read, which is an embarrassing thing for an evidence based book to admit and appears to be true.

Later work extended the idea into the question of why organisations struggle to learn from their own experience. The mechanism needs no patients and no operating theatres. It needs complexity, high stakes, and people who have to keep functioning.

A facilitated room has all three. And the person in that room carrying the most anxiety, most of the time, is the one standing at the front holding the pens.

The protocol:

The structure audit

Run this on your own session design, before the session.

List every structural element in the design. Every framework, every breakout, every template, every timed activity, every wall of sticky notes.

Against each one, answer a single question. What would happen if I removed this? Be specific. It is not enough to say that the session would be less organised. What would the room have to do that it currently does not have to do?

Mark every element where the honest answer is about you. It would be less predictable. I would not know where it was going. There would be silence. Somebody might say something I could not handle. The client would think I had not prepared properly.

Those are not bad reasons. They are real, and some of those structures should stay. The point is to know which ones are holding up the room and which ones are holding up you.

Then choose one and remove it. One element, in one session, where the reason was yours. Notice what the room does with the space.

In the room, when you feel the pull. There is a physical signal that comes just before you

reach for structure. For most people it sits in the chest or the hands, and it arrives around three seconds before the intervention. When you notice it, count to ten before you do anything. Most of what you were about to prevent will resolve itself in that time. What does not resolve is usually the most valuable thing available in that room.

Before your
next session

Which part of your standard design would you find hardest to run without, and is that because the room needs it?

4 What you did to their reasons



The energy was low after lunch, so something was added. A prize for the table with the most ideas. Points for each contribution. A competition between groups. A wall of stars. Perhaps it was your idea, and perhaps the client asked for it and would not be talked out of it.

The room became livelier immediately. Something else went quiet at the same moment, and it took a while to notice what.

Edward Deci found the same trade in 1971, using a puzzle.

What Deci found

Deci was a young psychologist at Carnegie Mellon University. He gave university students a puzzle called the Soma cube. It has seven oddly shaped blocks that fit together into a surprising number of forms. He asked the students to copy shapes from drawings, across three sessions on separate days.

In the middle of each session, Deci left the room for eight minutes. He said he would be back shortly. There were magazines on the table. Through a one way mirror, he timed how long each student kept playing with the puzzle when there was no longer anyone asking them to.

One group was paid a dollar for each puzzle they solved, and only during the middle session. The other group was never paid at all. In the final session, after the money had appeared and then stopped, the paid group spent less of their free time on the puzzle. The unpaid group's interest stayed where it was. Being paid, and then not being paid, had made the puzzle less worth doing than never being paid at all.

Deci explained this in terms of where we locate the reason for our own behaviour. Do something because you want to, and the reason sits inside you. Add a reward, and the reason can move outwards onto the reward. Take the reward away, and the original reason has already been pushed aside.

Two years later Mark Lepper, David Greene and Richard Nisbett took the idea to nursery children who plainly loved drawing. They promised some of the children a certificate for drawing, and gave it to them. Those children drew less afterwards than the children who had been promised nothing. They called it over-justification. Give people too many reasons to do something they already wanted to do, and the original reason weakens.

Every facilitator is asked to evoke group creativity. It sits in the competencies as C4, and the word is worth sitting with. To evoke

something is to call out what is already there. It is a very different instruction from manufacture it, and we have all spent an afternoon doing the second while reporting the first.

Where the evidence stops

This was argued over for two decades, and the argument is the useful part.

In the middle of the 1990s, Judy Cameron and David Pierce pooled the available studies and concluded that the harm was small, fragile and close to a myth. Deci, Richard Koestner and Richard Ryan reanalysed the evidence in 1999 across a large set of experiments, and the picture that came out was more specific than either side had claimed.

Praise and useful feedback tended to raise motivation. Unexpected rewards did very little in either direction. The rewards that reliably damaged interest were tangible, expected, and tied directly to the task. Do this, get that.

That is precisely the shape of most things a facilitator adds to a room to raise the energy.

The framework Deci later built with Ryan holds that a reward carries two messages at the same time. One says that you are good at this, which feeds motivation. The other says that you are doing this because somebody is making it worth your while, which takes motivation away. Which message lands depends on how the reward was set up, rather than on how large it was.

There is a cleaner demonstration outside the laboratory. Economists introduced a small fine for late collection at ten daycare centres in Haifa. Late collections went up. The fine had turned an obligation into a price, and parents who were willing to pay it stopped feeling bad.

The protocol:

The reason test

Run this on any session element designed to produce energy, participation or output.

Before you add it, ask what the reward is attached to. A reward tied to the act of taking part, such as most ideas, most dots, fastest table or first to finish, is attached to volume. Volume is what you will get. A reward tied to quality requires you to judge quality as it happens, which you usually cannot do. That is why almost every incentive in a room ends up attached to volume.

Check whether the reward was expected. Announced at the start, it becomes the reason for the work. Given afterwards without warning, it reads as recognition. This is the single largest difference in the research and it costs you nothing to apply. If you want to give something, give it at the end without having promised it.

Prefer specific feedback over anything tangible. The evidence is clear that verbal recognition tends to raise motivation while tangible rewards tend to lower it, and the effect depends on being specific. Saying that a particular reframing changed what the group was arguing about does real work. Saying great input does not.

When the client insists on the competition.

Two adjustments keep most of the value. Make the criterion something other than volume, and name that criterion before the activity begins. Then give the recognition for something the table could not have optimised for, such as the question the rest of the room had missed, or the objection that changed the design.

Watch what goes quiet. After introducing any incentive, notice what stopped happening. The tentative half formed ideas. The person who was thinking slowly. The objection that would have cost the table its lead. That is the price you paid, and it is usually paid by the most valuable contributions in the room.

Before your
next session

What are you adding to this design to create energy, and what is it teaching the room to produce?

5 The room that cannot hear it



Somebody said the important thing and it landed on nothing.

It is usually the newest person, or the one closest to the customer, or somebody from a function the rest of the room finds inconvenient. They made a short remark. The room acknowledged it politely, and moved on within about fifteen seconds. No one was rude. No one dismissed it. Six months later the same issue arrived as a serious problem, described as something that could not have been foreseen.

It was said, it was heard, and the room had no way to receive it.

Wesley Cohen and Daniel Levinthal gave that failure a name and a mechanism in 1990.

What Cohen and Levinthal found

Their paper appeared in a journal called Administrative Science Quarterly, and it started somewhere unexpected. In research on how memory works. New knowledge stays with us when it connects to something we already know. A production engineer who has spent years working with a particular machine will notice one line in a technical article that tells them something important. Somebody without that background reads the same line and sees nothing at all. The information was identical. The ability to receive it was not.

They called this absorptive capacity, and they argued that organisations work the same way. A company can only take in outside knowledge that connects to expertise it has already built. Everything else arrives as noise, however loudly it is said.

This builds on itself in an unpleasant direction. Invest steadily in understanding a market, and each new signal from that market becomes easier to recognise. Stop investing, and each new signal becomes harder to recognise, which makes the next one harder still. The ability to learn about something is built by having already learned about it.

Cohen and Levinthal also showed that this capacity, at the level of a whole organisation, depends on how knowledge is spread across people, how well those people talk to each other, and who sits at the boundary between the organisation and the outside world.

More recent work has turned this into a design question, by separating two distinct jobs. One job is going out and picking up the signal. The other job is acting on it. In most organisations those two jobs sit with different people, who report to different managers, and who meet in a room like yours for a few hours.

Which makes your room the place where the transfer either happens or fails.

The competencies describe this work under Guide Group to Appropriate and Useful Outcomes, and one component gets very close to it:

Help the group to make sense of underlying issues in their discussion

IAF Core Competencies, D2

Underlying is generous. Some issues are not underlying at all. They have been stated out loud, in plain language, by somebody sitting at the table, and the room has no way to attach them to anything it already knows. Our profession has a well developed vocabulary for what is hidden in a group. We have almost none for what is entirely visible and still cannot be received.

Where the evidence stops

Cohen and Levinthal studied whole firms. For many years the idea was measured using how much a company spent on research and development, which is a crude substitute and has been criticised heavily since.

There is no direct evidence connecting this mechanism to what happens in a single meeting. What you are getting in this chapter is a well supported organisational finding applied to a room by analogy. That analogy is mine and it is worth holding loosely.

It is worth holding at all because it explains something that otherwise looks like rudeness or politics, when it is neither. No one ignored the junior person. No one suppressed the customer's voice. The room simply had no prior knowledge to connect it to, so the remark arrived as unrelated, was treated with courtesy, and the session moved on.

The protocol:

Build the hook first

The design principle is that a room cannot absorb a signal it has nothing to attach it to. So you build the attachment point before the signal arrives.

In the design stage, identify who holds the outside knowledge. The person closest to the customer. The newest person. The one from the function everybody else finds inconvenient. Find out before the session what they know that the room does not.

Give the room the background first, rather than the conclusion. Twenty minutes before that person is going to speak, give the room the context that makes their contribution understandable. Not their point. The context their point depends on. If the customer's objection is about delivery reliability, the room needs to have spent twenty minutes on delivery before it can hear it.

Make one person responsible for translation, by name. Collapse the two jobs in your room into one. Name somebody at the start whose task is to restate anything that comes from outside the room in the room's own language. Their job is not to agree with it. Their job is to make it understandable.

Use the unattached remark test. When something is said and the room responds with brief courtesy and moves on, stop. Say that

you want to return to what was just said, and ask the room one question. What would have to be true for that to be the most important thing said today? This forces the room to build the connection it could not build on its own.

Capture what the room could not use. At the close, list anything that was raised and that the session had no way to absorb. Give that list to whoever commissioned the session, and say clearly what it is. You will meet these items again in six months, described as a surprise.

About Flyntrok

We work with organisations going through change. Sometimes that is growth, sometimes a shift in direction, sometimes a rebuilding of how people work together. We are usually brought in when something has to move, and the people who have to move it are the ones who will live with the result.

We are interested in research, and we work with it every day. Evidence is not something we consult before a large engagement and then put away. It sits underneath how we design sessions, how we advise, and how we argue with each other about what will actually work in a room. When we are unsure, we go and read. When the research contradicts what we have been doing, we change what we have been doing, which happens more often than we would like.

It also means we say we do not know rather more often than is commercially convenient. A good deal of what is sold as certainty in this field is somebody's preference with a diagram attached.

The name Flyntrok comes from the flint used to make fire. Sparked with care, on purpose, by the people who need the fire.

Every week we publish a piece in this series, On the Evidence. Each one takes a body of

research and works out, in plain language, what it means for people who convene, lead and facilitate. Some of those pieces are in this book. Most of them are not.



**Scan to read more from the
On the Evidence series, and
everything else we publish.**

One question before you go

If one of these chapters made you want to change something on Monday, we would like to know which one.

This is partly research and partly self interest. There are more chapters written than could fit in a book this size, and no clear sense of which five were the right five. Your answer settles an argument we are currently having with ourselves.

Write in at programs@flyntrok.com

FLYNTROK